Investigating the dynamic role of fluctuations in ongoing activity in the human brain

Dissertation submitted for the degree of Doctor of Philosophy of the University College London

Maren Urner

October 2013

Institute of Cognitive Neuroscience
Wellcome Trust Centre for Neuroimaging
Institute of Neurology
University College London

Declaration

I, Maren Urner, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

The work presented in Chapters 3 and 4 has been published as the following papers:

Urner, M., Schwarzkopf, D. S., Friston, K., Rees, G. (2013). Early visual learning induces long-lasting connectivity changes during rest in the human brain. Neuroimage 77, 148-56.

Urner, M. Sarri, M., Grahn, J., Manly, T., Rees, G., Friston, K. (2013). The role of prestimulus activity in visual extinction. Neuropsychologia 51(8), 1630-7.

"The fact that the body is lying down is no reason for supposing that the mind is at peace. Rest is… far from restful."
Seneca (~60 A.D. (1969))

Abstract

Traditionally, the focus in cognitive neuroscience has been on so-called evoked neural activity in response to certain stimuli or experiences. However, most of the brain's activity is actually spontaneous and therefore not ascribed to the processing of a certain task or stimulus – or in other words, uncoupled to overt stimuli or motor outputs. In this thesis I investigated the functional role of spontaneous activity with a focus on its role in contextual changes ranging from recent experiences of individuals to trial-by-trial variability in a certain task. I studied the nature of ongoing activity from two perspectives: One looking at *changes* in the ongoing activity due to learning, and the other one looking at the *predictive role* of prestimulus activity using different methodologies, i.e. EEG and fMRI. Finally, I ventured into the realm of inter-individual differences and mind-wandering to investigate the relationship between ongoing activity, certain behavioural traits and neuronal connectivity.

Contents

List of Figures

List of Tables

Chapter 1 General introduction

1.1 Spontaneous and evoked activity

Traditionally, the focus in cognitive neuroscience has been on so-called evoked neural activity in response to certain stimuli or experiences. However, most of the brain's activity is actually spontaneous and therefore not ascribed to the processing of a certain task or stimulus – or in other words, uncoupled to overt stimuli or motor outputs. Possibly, the existence of ongoing intrinsic activity was first noted by Hans Berger when he introduced electroencephalography for humans in 1929 (Berger, 1929), asking whether "it [is] possible to demonstrate the influence of intellectual work upon the human electroencephalogram, insofar as it has been reported here?" to conclude subsequently that "[o]f course, one should not at first entertain too high hopes with regard to this, because mental work, as I explained elsewhere, adds only a small increment to the cortical work which is going on continuously and not only in the waking state". Four years later, Bishop (1933) reported the potential physiological significance of the ongoing activity describing his experiments with rabbits. He observed cyclic changes in the excitability in visual cortex during stimulation of the optic nerve. Summarising his findings, he stated that "[…] we would look upon the cortex as being in constant activity, the physiological activity of the whole network of neurons bearing some direct relationship to the 'present state' of the animal's complex behavio[u]r which is sometimes referred to as his 'mental state'".

Indeed, ongoing activity occurs throughout the brain and its existence is manifested in the variability of cortical responses in repeated responses to physically identical

conditions or stimuli. In the past, this variability had simply been labelled as noise and scientists got rid of it by averaging over repeated trials (Gerstein, 1960; Zohary et al., 1994). However, during the last two decades an increasing number of neuroscientists recognised that ongoing neural activity is not mere noise, but plays a fundamental role in stimulus-driven processing (Arieli et al., 1996; Tsodyks et al., 1999) and behavioural variability indeed (Hesselmann, Kell, Eger, et al., 2008; Coste et al., 2011; Kleinschmidt et al., 2012).

I investigated the characteristics of the ongoing brain activity[1] focusing on its functional role and its role in contextual changes, where contextual changes can be differences in the experience of individuals (e.g. learning-related changes) or can be related to trial-by-trial variability.

1.2 The study of spontaneous activity

Why study ongoing brain activity? Contrary to the focus on evoked activity in neuroscience, spontaneous neural activity dominates the brain's energy consumption (Attwell and Laughlin, 2001; Mintun et al., 2001; Attwell and Iadecola, 2002). The energy consumption during rest exceeds task-related increases in neural metabolism, which are usually < 5 % (Raichle and Mintun, 2006). Thus, the majority of neuroscientific studies are focused on a minor component of brain activity. Maybe it is time for an adjustment or alteration in the neurosciences, shifting towards an

[1] In the literature, different terms have been used to describe ongoing neural activity – as compared to evoked responses – among which are "resting state activity", "endogenous activity", "spontaneous activity", and "autonomous activity". I use the term "ongoing activity" and "spontaneous fluctuations" interchangeably and refer to activity not evoked by an external stimulus or task.

experimental approach that is indeed focusing on the factor that uses the lion's share of the brain's energy, namely ongoing or spontaneous neural activity.

Although ongoing brain activity has been studied using electrophysiological and neuroimaging methods, its physiological origin and cognitive consequences are not yet fully understood. Crucially, any clarification is difficult by its very nature, because any study that addresses the functional significance of spontaneous fluctuations inevitably requires a primary task-context in order to probe perceptual and / or behavioural consequences of the fluctuations (Hesselmann, Kell, and Kleinschmidt, 2008). Attributed roles of ongoing brain activity span processes at different levels of neural activity and range from the traditional view of "intrinsic noise" over low-level physiological processes and uncontrolled mental activity to a monitoring of the environment (Mantini and Vanduffel, 2013). In conclusion, one of the most intriguing questions in the neurosciences might be related to the functional significance of the brain's "intrinsic noise".

1.2.1 Electrophysiological research of ongoing activity

The brain is a noisy system whose processing parts – the neurons – receive a large number of fluctuating inputs which in turn generate spike patterns. These often appear very irregular and much of the activity is spontaneous.

1.2.1.1 Cortical states and response variability

Cortical states are determined by the states of individual neurons and the states of individual neurons are in turn related to the state of their neighbours. Possibly the ground-breaking study investigating spontaneous activity and its relation to the large variability of evoked responses to repeated presentations of the same stimulus, is the

work by Arieli et al. (1996). They showed that single trial responses in cat visual cortex can be predicted by the linear summation of the deterministic response and the preceding ongoing activity, concluding that ongoing activity plays a noteworthy role in cortical function and for cognitive processes. Given the observation that ongoing fluctuations influence cortical processing, one might wonder how these two types of brain processes interact. In order to answer this question, the relation between the activity of single neurons and the dynamics of the network in which they are embedded has been explored. Using single-unit recordings and real-time imaging, Tsodyks et al. (1999) showed that the firing rate of a spontaneously active single neocortical neuron depends on the instantaneous spatial pattern of ongoing population activity of a larger cortical area. The spatial patterns of population activity recorded during spontaneous firing and those when driven by the optimal input were very similar. Moreover, the correspondence between evoked neural activity and the structure of an input signal seems to mature with age due to a shift in the dynamics of spontaneous activity (Fiser et al., 2004). These results suggest that evoked neural activity in response to sensory stimulation might represent modulated and triggered ongoing neural circuit dynamics, instead of the structure of the input itself. On a more theoretical level, aspects of neuronal responses that are often considered as "noise" might be essential components of the way in which information is propagated or represented in neurons (Ermentrout et al., 2008).

Furthermore, the spatio-temporal structure of spontaneous activity has been shown to be highly coherent over different networks and in multiple species (Chiu and Weliky, 2001). This coherent spontaneous activity of neurons involves a set of dynamically

switching cortical states similar to the well-known orientation maps (Kenet et al., 2003).

Taken together, the traditional belief that neural activity is primarily driven by sensory input from the environment might be out-dated. Instead, spontaneous activity of single neurons and their networks seems to play a crucial role for the cortical processing following a sensory stimulus. Correspondingly, this observation might explain the well-established response variability observed after the repeated presentation of physically identical stimuli (Henry et al., 1973; Vogels et al., 1989; Azouz and Gray, 2008). The study of response variability is an important tool used to examine the role of spontaneous activity related to the question of whether it carries predictions about sensory stimuli. Specifically, the investigation of so-called prestimulus activity occurring prior to stimulus presentation has become a standard approach, based on the assumption that predictions – if any – are expressed shortly before stimulus onset.

However, the linear relationship between spontaneous and evoked activity as proposed by some studies (Arieli et al., 1996; Azouz and Gray, 2008) has been challenged using neuroimaging methods (e.g. Hesselmann, Kell, & Kleinschmidt, 2008; Schölvinck et al., 2012) and might indeed differ across the brain. Recently, it has been shown that evoked activity across a large part of the human cortex interacts negatively with ongoing activity (He, 2013). Consequently, a higher prestimulus baseline resulted in less activation – or more deactivation – leading to a decreased trial-to-trial variability of cortical activity after stimulus onset. Thus, measuring

across-trial variability might provide a new approach to interpret neuroimaging experiments using trial-based data.

1.2.1.2 Predictive coding and predictive timing

Independent of how spontaneous activity interacts with evoked responses, its existence and functional influence as described above begs the question of why it is there in the first place. One possible explanation for its existence is the concept of predictive coding. The basic idea is that the brain possesses internalised representations of the world which it uses to "predict" what happens in the environment thereby inferring the most likely cause of sensory events (Friston, 2005). In other words, the brain might generate hypotheses about the potential causes of upcoming sensory events and compares these hypotheses with incoming sensory information (Summerfield et al., 2006). The difference between the two – the so-called prediction error – is then propagated forward throughout the cortical hierarchy; internal representations might be adjusted subsequently. Lately, the theory around predictive coding has been extended by the notion of predictive timing (Schroeder and Lakatos, 2009; Arnal and Giraud, 2012), which exploits the temporal regularities or associative contingencies (e.g. the temporal relation between two inputs) to infer the occurrence of forthcoming sensory events. In short, predictive timing is the process by which uncertainty about the temporal occurrence of events is minimised such that their processing and detection are facilitated (Jones et al., 2002; Nobre et al., 2007). During the last decade, both accounts – predictive coding and predictive timing – have been supported by several studies providing evidence for the idea that the hypotheses – or predictions – generated by the brain might be embodied in the spatial and temporal structure of spontaneous activity.

One proposed neurophysiological substrate of the operations described in the framework of predictive coding (i.e. directional message-passing) are cortical oscillations, i.e. those observed as ongoing activity. Both, gamma and beta oscillations (see 2.4 for a description of frequency bands) have been shown to play a role in predictive coding. For instance, gamma activity scales with prediction errors (Arnal et al., 2011) and has been implicated in the evaluation of sensory predictions, possibly depending on the match between bottom-up input and expectations (Herrmann et al., 2004). Beta activity has been associated with error-related effects as well, but in a different direction of processing than gamma, i.e. downstream from prediction error generation (Fujioka et al., 2009; Iversen et al., 2009). Combining these findings, gamma and beta could underlie the information flow in opposite directions, i.e. forward versus backward (for a review see Wang, 2010). Thus, prediction errors might be transmitted in a feed-forward manner using the gamma frequency channel, while predictions and their reconsiderations could be propagated by the beta channel in a backward direction (Chen et al., 2009; Wang, 2010).

With regard to predictive timing, low and mid-frequency oscillations have been identified as important. Their interactions during temporal expectations support a functional cooperation between these oscillations in predictive timing (Saleh et al., 2010). For instance, the anticipation of sensory events resets the phase of slow, delta-theta activity before stimulus presentation and the predictive alignment of these oscillations (in an optimised excitability phase) results in faster stimulus detection (Lakatos et al., 2008). Also, reduced early sensory responses are observed in response to stimuli that are implicitly expected based on their temporal regularity (Costa-Faidella et al., 2011). Furthermore, neural and perceptual responses for

temporally unpredictable stimuli are modulated by the phase of alpha oscillations at which stimulation occurs (Busch et al., 2009). However, how the oscillations at the different low frequency bands interact during predictive timing of sensory events remains to be solved.

In summary, there are several lines of electrophysiological research providing evidence which supports the idea that spontaneous activity has certain predictive power for subsequent neural responses to upcoming sensory stimuli.

1.2.2 Neuroimaging research of ongoing activity

The crucial difference between electrophysiological methods and neuroimaging using functional MRI lies in the temporal resolution: while dynamically switching cortical states can be accessed with single cell recordings for example, functional MRI is restricted to a much broader timescale. Nevertheless, the technique has been used to study spontaneous activity for almost twenty years based on the pioneering work of Biswal et al. (1995), who first observed correlations of low frequency fluctuations in the resting brain between different regions involved in a simple motor task. He concluded that these might be a manifestation of functional connectivity in the brain, contrary to the traditional view of the brain as being driven by transient environmental demands. Similar to other neuroscientific methods, functional neuroimaging has spent most of its infancy with studies of evoked responses to sensory, cognitive and motor events (Posner and Raichle, 1994). Therefore, it took several years before the examination of ongoing activity became a major topic among neuroscientists using functional MRI.

1.2.2.1 Resting state fluctuations

The presence of low-frequency fluctuations in the blood oxygen level dependent (BOLD) signal (see 2.3.2) of functional MRI is a well-established finding in neuroimaging. Recent studies have identified these fluctuations as a potentially important manifestation of spontaneous neuronal activity, which is indeed organised in distinct patterns – often referred to as resting state networks (De Luca et al., 2005; Fox and Raichle, 2007; Corbetta, 2012), which will be described below (see 1.2.2.4). Several observations made about resting state fluctuations appear to become known principles. First, their spatial organisation is preserved in humans under anaesthesia (Greicius et al., 2008) and during early stages of sleep (Larson-Prior et al., 2009), as well as in monkeys (Vincent et al., 2007) and rats (Lu et al., 2007). More than that, the fluctuations in ongoing activity show a high consistency and reproducibility across participants and sessions (Damoiseaux et al., 2006; Chen et al., 2008; Meindl et al., 2010; Wang et al., 2010), a high test-retest reliability (Shehzad et al., 2009; Wang et al., 2010) and a high reproducibility across different analytic approaches (Long et al., 2008; Franco et al., 2009). Second, the strength of coherence between different regions exhibiting ongoing activity varies with different parameters among which are age (Fair et al., 2008), experience (Lewis et al., 2009) and disease (Zhang and Raichle, 2010). For instance, it has been argued that disruption in the resting state coherence between different nodes might be a sensitive early biomarker for certain diseases, such as Alzheimer's (Zhang and Raichle, 2010). Third, the slow global fluctuations that are observed in the BOLD signal at rest can be of the same magnitude as signal changes in response to task-related paradigms (Damoiseaux et al., 2006). Fourth – related to the previous aspect and extending the

electrophysiological findings discussed above in 1.2.1 – they contribute significantly to the variability in evoked signals (Fox et al., 2006) and the variability of the associated behavioural response (Fox et al., 2007). Last – but certainly not least – the frequency distribution of the spontaneous fluctuations is significantly different from BOLD fluctuations observed in water phantoms (Zarahn et al., 1997).

Taken together, the aforementioned studies and findings have provided insight into the intrinsic functional architecture of the brain. However, the mechanisms underlying these resting-state fluctuations are still controversial.

1.2.2.2 Vascular basis

A possible explanation for the observed global anatomy of spontaneous activity – given its reproducibility – which is not linked to neural architecture, could be based on vascular mechanisms. All voxels in the brain show a global level of coherence between each other. Thus, the origin of the observed spontaneous fluctuations might lie in changes of blood flow. One possible explanation is linked to so-called draining veins, which exhibit a form of "blood stealing" throughout the tissue whereby active regions generate blood flow increases at the expense of other nearby regions. Alternatively, further poorly understood mechanisms of vascular regulation could play a role (Buckner et al., 2008). Indeed, changes in respiration and intracranial pressure are known to influence blood flow and blood oxygenation in the brain; the BOLD signal measured during functional MRI is based on hemodynamic measures of blood flow and only indirectly linked to neural activity (see 2.3.2). Variations in heart rate and respiration are known to contribute to fMRI resting-state fluctuations (Wise et al., 2004; Birn et al., 2006). The highest power of spontaneous fluctuations

is below 0.1 Hz (Cordes et al., 2001), which is beneath normal breathing frequency. Nevertheless, the aforementioned variations in respiration rate and depth are observed at these low frequencies (Birn et al., 2006) and correlate with the BOLD signal acquired during rest. In particular, these correlations overlap with the areas characterised as the default mode network (see 1.2.2.4 for a detailed description of this particular resting state network) (Birn et al., 2006).

Finally, with regard to the potential role of vascular coupling for spontaneous fluctuations, two aspects are worth mentioning. First, resting state networks can also be identified using measures of resting glucose metabolism, i.e. without relying on vascular coupling at all (Vogt et al., 2006). Second, confounds in resting-state data due to cardiac and respiratory effects are now addressed by sophisticated methods (Glover et al., 2000; Birn et al., 2008), which will be described in 2.3.4.

1.2.2.3 Neural basis

Knowing that part of the spontaneous fluctuations depend on non-neuronal physiological factors, it needs to be clarified what the *neural* basis of the signals is. Indeed, an effort has been made to determine the electrical correlates of the fMRI BOLD signal (see Khader, Schicke, Röder, & Rösler, 2008; Logothetis, 2008 for summaries of this work involving different perspectives). The conclusion drawn from this work is that the BOLD signal is best correlated with local field potentials (LFPs). Probably the most cited study in this line of research observed that the power of LFPs recorded in monkey visual cortex fluctuates at a similar timescale to the one of spontaneous fluctuations measured by fMRI, i.e. the previously mentioned < 0.1 Hz (Leopold et al., 2003). Similar to the long-distance coherences observed in

resting-state activity, the fluctuations observed by Leopold et al. (2003) show a high coherence across electrode pairs without any changes due to cortical distance. Comparable fluctuations have been observed using intracranial electrocorticography (ECoG, i.e. surface electrodes) recordings in humans (He et al., 2008; Nir et al., 2008). Importantly, recordings of neuronal electrical activity either from the scalp with electroencephalography (EEG) or from the surface of the brain with ECoG always encompass a summation of a population of LFPs. One possibility to analyse fMRI resting-state fluctuations and its underlying neural activity is the simultaneous recording of EEG and fMRI (Goldman et al., 2002; Laufs et al., 2003; Mantini et al., 2007). Conventionally, LFPs are described according to their frequency components and two of these components have been associated with spontaneous BOLD fluctuations: First, fluctuations in the power of higher frequencies (Leopold et al., 2003), which are associated with cognition (Donner et al., 2009; Uhlhaas et al., 2010), and second, fluctuations that approximate those present in the spontaneous BOLD signal, commonly summarised as slow cortical potentials (SCPs) (Rockstroh et al., 1989; He et al., 2008; Raichle, 2011). Increasing evidence supports the idea that spontaneous BOLD fluctuations are most correlated with LFP activity in the SCP range (Lu et al., 2007; He et al., 2008; He and Raichle, 2009), while the spatial patterns of coherence are maintained across different levels of consciousness ranging from wakefulness to rapid eye movement (REM) sleep and slow wave sleep (He et al., 2008). For instance, alpha power (see 2.4 for a description of frequency bands) is negatively correlated with spontaneous BOLD fluctuations in occipital cortex, as well as in inferior frontal and superior parietal regions (Goldman et al., 2002; Laufs et al., 2003; Moosmann et al., 2003). The latter regions are commonly active during

rest and are known to be involved in the default mode network (see 1.2.2.4 for a detailed description of this network). The less persistent correlation of power in the higher frequencies (i.e. the spatial correlation with the BOLD signal is only observed during wakefulness and REM sleep) is in line with the known role of these frequencies for consciousness awareness (Fries, 2009).

One way to analyse simultaneously acquired EEG and fMRI data is to use so-called independent component analysis (ICA). It is used for BOLD data in order to identify different networks and to correlate these subsequently with different frequency bands of the EEG data. In doing so, Mantini et al. (2007) found multiple resting-state networks in the BOLD data, which correlated with different frequency bands of the EEG data.

Finally, Shmuel and Leopold (2008) used intracortical neurophysiological recordings in combination with fMRI to investigate the relationship between spontaneous fMRI and LFP signals as directly as possible. They found strong correlations between the spiking rates (of the neurons close to the recording electrode) and slow fluctuations in the fMRI signal, as well as with slow power changes in the multi-unit activity (> 100 Hz) and LFP band of higher frequencies (25 – 80 Hz).

In summary, a large number of studies using different methodological approaches including simultaneous measurements of fMRI and neuronal activity have established the presence of a coupling between slow fluctuations in the BOLD signal measured with fMRI and underlying fluctuations in the neural activity across multiple regions, frequencies and states of consciousness. Thus, modulations in both signals possibly share the same origin – which might be subcortical in nature. Future

studies are required to determine the origin the slow fluctuations in spontaneous activity.

1.2.2.4 Functional networks

Independent of the underlying mechanisms of spontaneous brain activity, the organisation of the commonly observed low-frequency fluctuations of spontaneous activity organised into specific and distinct patterns is a robust finding. A number of so-called resting state networks has been discovered since the pioneering study of Biswal et al. (1995). Thus, the large-scale sensory-motor network reported back then was only the first one of multiple networks observed during rest and exhibiting a high similarity to task-activated networks. Comparable relationships have been found for other modalities, such as visual, auditory and language processing networks (Lowe et al., 1998; Hampson et al., 2002; van de Ven et al., 2004). For instance, regions in intraparietal sulcus, frontal eye field, and middle temporal cortex that are commonly recruited during visuospatial attention or oculomotor tasks also form a functional network at rest different from the classical "visual" network including V1 to V4. Probably the most famous of the these networks is the so-called default mode network (DMN), which had been first described by Raichle et al. (2001) and which is most consistently found across experiments (Anticevic et al., 2012; Mantini and Vanduffel, 2013). Raichle's et al. (2001) idea of a baseline – or default state of the brain – stems from a common finding in positron emission tomography (PET) and functional MRI studies, i.e. the observation that a certain set of brain regions shows decreases in activity independent of a particular task and with little variation in their location across a wide range of tasks (Shulman et al., 1997). Shulman et al. (1997) provided the first formal characterisation of task-induced activity decreases. Based

on their finding, the idea of a default state of brain function distinct from the sates involved in sensory and executive functions (i.e. in attention-demanding, goal-directed behaviours) has been established (Buckner et al., 2008). By use of a generally accepted quantitative circulatory and metabolic definition of brain activation, the authors established criteria for a baseline state characterised by the absence of activation. In the meantime, the default mode network has been studied extensively, see Buckner et al. (2008) for a review. Anatomically, it is considered to encompass certain key regions among which are posterior cingulate cortex (PCC), left and right inferior parietal cortex, and ventromedial prefrontal cortex (vmPFC). In addition, lateral temporal cortex, dorsal medial prefrontal cortex (dmPFC) and the hippocampal formation have been associated with the default mode network (Buckner et al., 2008).

In general, resting state patterns of coherence characterised as "resting state networks" take patterns of anatomical connectivity in the human (Zhang et al., 2008) and the monkey brain (Vincent et al., 2007) into account, but are not restricted to these anatomical connections. Therefore, the absence of monosynaptic connections between certain brain regions does not exclude the existence of functional connectivity as for instance expressed in the networks of resting state coherence (Raichle, 2011).

1.3 The functional role of spontaneous activity

Having discussed the existence and potential neurophysiological basis of spontaneous fluctuations and their networks observed during rest, I will now focus on their functional role. As mentioned earlier, traditionally, ongoing activity has been

often considered as “intrinsic noise” that can be averaged over trials in order to get to the essence of the “real signal”. Even in view of the structural regularities, i.e. the organisation into networks, it could still be possible that intrinsic functional connectivity merely reflects some “noise” that happens to have a non-random structural connectivity carrying the described form of spatial patterns. In support of this notion, computational and empirical studies have shown a correspondence between intrinsic functional and anatomical connectivity (Sporns et al., 2000; Greicius et al., 2009). However, at the level of the entire brain and involving a systematic quantitative analysis, the match is not impeccable. Honey et al. (2009) found that structural connectivity could predict the empirically observed functional connectivity only to a limited extend. Leaving potential limitations during data collection and analysis – which could cause the deficiencies of predicting functional from structural connectivity – aside, an alternative hypothesis has been put forward: If structural connectivity merely forms the backbone of functional connectivity, it has to be influenced and shaped by additional context-dependant modulations (Sadaghiani, Hesselmann, et al., 2010).

Although this hypothesis might be contradictory given the high consistency of spatial patterns of ongoing activity across different levels of consciousness and context as mentioned in 1.2.2.3, several lines of research have accumulated evidence supporting the hypothesis that ongoing brain activity is indeed context sensitive. Thus, even though functional connectivity patterns are persevered qualitatively in a range of states, including light and deep sleep (Horovitz et al., 2008, 2009; Nir et al., 2008), as well as severe disorders of consciousness such as vegetative state (Boly et al., 2009), they do express quantitative changes. For example, functional connectivity in

the DMN decreases with the degree of consciousness in healthy individuals (i.e. between wakefulness, deep sleep and a state of physiological unconsciousness (Horovitz et al., 2009), as well as in patients (i.e. between minimally conscious state, vegetative state and coma) (Vanhaudenhuyse et al., 2010). As a last counter-argument to the idea of ongoing activity as mere "noise", I refer to the finding that the reduction in connectivity between regions of the DMN, i.e. frontal and posterior areas, during sleep is anatomically selective. Therefore, it is very unlikely that intrinsic connectivity only reflects a change in the level of "noise" that is propagated through an anatomically connected network (Sadaghiani, Scheeringa, et al., 2010).

In the following, I focus on the different lines of research supporting the idea that ongoing activity is context-sensitive and also can influence behaviour.

1.3.1 Perceptual domain

The perceptual impact of prestimulus activity fluctuations has been investigated in different perceptual paradigms. For example, Boly et al. (2007) used a somatosensory detection task and found that prestimulus activity in a large distributed network determined whether stimuli close to perceptual threshold were detected or not on a single trial basis. Regions involved in the predictive system included the thalamus, dorsal anterior cingulate cortex (dACC), parieto-frontal areas and inferior frontal gyrus. These are all regions that are commonly active in response to a wide range of (cognitive) tasks (Smith et al., 2009a; Corbetta, 2012). In contrast, prestimulus activity in regions that have been coined as "task-negative" belonging to the DMN, i.e. posterior cingulate (PCC), parahippocampal and lateral parietal components, was higher for trials that were missed by participants. These results

support the concept of a simple dichotomy between "task-positive" and "task-negative" networks whose activity either facilitate or deteriorate perception. However, more recent findings revealed a more complicated picture highlighting that context plays an important role predicting whether the ongoing activity in a certain brain area plays a role for subsequent perception of a stimulus – or not. Sadaghiani et al. (2009) used a free-response, auditory detection task and examined whether prestimulus activity was predictive of when participants perceived auditory stimuli at individual detection threshold. A rather complex picture was observed: comparing detections and misses, the former ones were preceded by higher activity in early auditory cortex as well as in a network including thalamus, anterior insula and dACC; misses were preceded by higher activity in the so-called dorsal attention system, including frontal and parietal areas. Thus, two task-positive networks showed opposite effects. In addition, regions of the DMN were more active prior to hits, not misses. The time courses of prestimulus effects in the different networks were distinct, underlining the idea that the two observations were probably not mere mirrored effects. These results highlight the complexity of ongoing activity – which is not only organised in spatially defined networks. Mora than that, these networks are independently organised and context sensitive. In other words, the context – including specific sensory features and cognitive faculties – of a perceptual decision impacts to what extent ongoing activity in different networks determines stimulus processing and eventually perception.

In line with this observation, several researchers have speculated that a perceptual task involving choices – compared to all-or-none detection tasks – might affect prestimulus activity in a more localised way, i.e. restricted to a certain area that is

task-relevant compared to more general brain networks related to attention or memory. Two linked fMRI studies have provided evidence for this idea. First, increased prestimulus activity in fusiform face area (FFA), a region that has been previously related to the processing of faces, was observed when perceiving a face compared to a vase in a paradigm using Rubin's ambiguous vase-face figure (Rubin, 1915; Hesselmann, Kell, Eger, et al., 2008). Second, prestimulus activity in human MT+, an area crucial for motion processing located in the middle temporal cortex, was higher in trials when participants detected coherence in a random dot stimulus paradigm compared to trials when they did not (Hesselmann, Kell, and Kleinschmidt, 2008). The difference was found comparing trials that were physically identical, i.e. at a coherence level that resembled the individual detection threshold of coherence (see Chapter 5 for a more detailed description of these findings that were used as the basis of the study described there).

In addition to the studies using fMRI, several EEG and MEG studies have provided additional support for the functional role of ongoing brain activity for perception. Suffering from poorer spatial resolution, these studies provide the benefit of identifying specific oscillations bands. For instance, using MEG, visual discriminability has been shown to decrease with an increase of certain low oscillations, i.e. in the so-called alpha band (see 2.4 for a more detailed description of frequency bands), in occipital-parietal cortex (van Dijk et al., 2008). Comparable to the findings in the fMRI literature, a free-response task revealed a rather complicated picture of different frequency bands being important for perceptual outcome. Using EEG and a somatosensory threshold detection task, Linkenkaer-Hansen et al. (2004) found that prestimulus fluctuations at medium power levels of several frequency

bands over somatosensory cortex resulted in highest detection rates and shortest reaction times. In contrast, over parietal electrodes the best behavioural performance was associated with the highest power in the same frequency bands. In addition to power analyses of oscillations, the phase of slow oscillations has been shown to have certain predictive power for perceptual outcome, e.g. in visual threshold detection tasks (Thut et al., 2006; Busch et al., 2009; Mathewson et al., 2009). Thus, the cortex might go through different states of excitability – so-called microstates – which can differ in speed, depending on the oscillation whose phase is important (Monto et al., 2008; Busch et al., 2009; Mathewson et al., 2009).

The predictive nature of prestimulus activity with regard to perception has also been shown looking at electrode recordings in monkeys (Supèr et al., 2003). Recording from primary visual cortex, it has been observed that reported stimuli were preceded by higher and more correlated neural activity compared to not-reported ones. The authors concluded that the strength of neural activity and the functional connectivity between different neurons in primary sensory areas contributes to perceptual processing. More precisely, visual cortex needs to be in a suitable state to result in subsequent stimulus detection.

In conclusion, numerous lines of research using different techniques and methodological approaches have presented supporting evidence to the idea that variability in perceptual performance can be – partly – explained by the variability in intrinsic – ongoing – processes in the brain; different signals measuring brain activity directly or indirectly can be used to forecast perception in the human and primate brain.

1.3.2 Motor domain

The functional role of ongoing activity has also been investigated with regard to motor activity and behaviour. In a series of two studies, Fox et al. (2007, 2006) showed first that trial-to-trial variability of finger movement-related activity in motor cortex can be largely attributed to fluctuations in ongoing activity (Fox et al., 2006), and second, that variability in behaviour depends on spontaneous activity as well (Fox et al., 2007). In order to do so, they used a simple button press task with the right index finger and a trick in order to disentangle evoked from ongoing activity. They exploited the fact that right and left somatomotor cortex exhibit correlated spontaneous activity (Biswal et al., 1995; Cordes et al., 2000) and that button presses with one hand do not result in evoked responses in ipsilateral motor cortex. More precisely, they used activity in the right motor cortex as a proxy of spontaneous activity in the left motor cortex (activated by the right hand button presses). In doing so, they showed that 40 % of the trial-to-trial variability in the BOLD response in left motor cortex can be ascribed to spontaneous activity. Based on the additional observation that the trial-by-trial evoked activity did not depend on whether the spontaneous activity in a given trial was high or low, they concluded that both types of activity are superimposed in a linear way. The second study made use of the same paradigm, but distinguished button presses according to their strength, i.e. trials were rated as either soft or hard button presses. Subtracting right motor cortex estimates of spontaneous activity from the activity measured in left motor cortex eliminated the apparent difference between responses to soft and hard button presses. In doing so, the study showed that 74 % of the relationship between spontaneous force variability and BOLD activity in left motor cortex can be explained by spontaneous activity

fluctuations. In sum, spontaneous activity influences trial-to-trial variability on the neuronal and behaviour level in response to a simple motor task.

On a different level, Ramot et al. (2011) confirmed a link between resting state activity and spontaneously emerging subconscious oculomotor behaviour. The behaviour they looked at, is a type of eye movement that occurs spontaneously and subliminally whenever we close our eyes (Allik et al., 1981; Collewijn et al., 1985). However, the neuronal mechanisms and functionality of these spontaneous eye movements are largely unknown. The findings that spontaneous fluctuations in the BOLD signal were correlated to the amplitude and velocity of these eye movements and that the neuronal activity was linked to coordinated motor programs (involving oculomotor neurons and muscles), provide further evidence for the idea that neuronal activity related to movement and associated behaviour is influenced by spontaneous activity in the brain.

1.3.3 Cognitive domain

The distinction between perceptual, motor and cognitive tasks is not straight-forward, because usually experimental tasks involve all three domains to a certain and varying degree. Often, the cognitive level is considered as the "highest" one, involving specific brain regions, so-called higher cognitive brain areas compared to primary sensory and motor areas. Given the difficulty of separating domains, I will introduce this paragraph with some studies bridging the gap between perception, movement and cognition using inhibitory control. Inhibitory control refers to the ability to suppress planned or ongoing processes, which might be related to movements or cognition. Using a monotonous task and MEG, Mazaheri et al. (2009)

showed that false alarms were preceded by higher alpha band power in visual areas and bilateral somatosensory cortices compared to correct withholds. An EEG study that looked at single-trial EEG topographic analyses to avoid averaging out effects that might get lost in the more typical ERP analyses (see 2.4.1), found supporting evidence for the idea that fluctuations in the ongoing activity immediately preceding stimulus presentation contribute to behavioural outcomes in an inhibitory control task (Chavan et al., 2013). They used an auditory spatial go no-go task and observed that a specific configuration of the EEG voltage field manifested itself before correct rejections, but not before false alarms. Using source estimation on the EEG topography, a fronto-parietal network was identified. These results support the idea that prestimulus brain activity also influences behavioural outcomes in an inhibitory control task. Furthermore, the identification and involvement of the fronto-parietal network suggests that the state-dependency of sensory-cognitive processing includes high-order, top-down inhibitory control mechanisms.

Until now, there are only a few "purely" cognitive control studies investigating the relation between ongoing brain activity fluctuations and inter-trial variability in evoked responses. However, the same picture emerges, confirming the crucial role of ongoing activity for evoked neural responses as well as behavioural outcome. For instance, Coste et al. (2011) used the Stroop task with colour-word interference and fMRI to show that prestimulus activity in several task-relevant brain regions (including dorsal anterior cingulate, dorsolateral prefrontal cortex and ventral visual areas) predicted subsequent response speed. Furthermore, this effect scaled with the Stroop effect size, i.e. it was only significant in participants who exhibited behavioural interference.

Another approach in the realm of the study of ongoing activity is focused on the relation between clinical phenomenon and changed patterns in resting state connectivity. For instance, a recent review identified 16 fMRI studies that investigated the use of resting-state fMRI in major depression (MD) and concludes that this research has yielded insight into the pathophysiology of depressive symptoms. Foremost, the role of the cortico-limbic mood regulating circuit and the interaction between task-positive and -negative networks in MD are emphasised (L et al., 2012).

Yet another methodological approach looks at the functional role of intrinsic connectivity on cognition by examining specific patterns and their capability to change in response to certain cognitive experiences. As soon as after one scanning session using fMRI, i.e. a time span that most likely does not include any gross anatomical changes, intrinsic functional connectivity has been reported to be sensitive to visuo-motor learning (Albert, Robertson, and Miall, 2009), episodic memory (Tambini et al., 2010), as well as language tasks (Waites et al., 2005; Hasson et al., 2009). To conclude, the functional context of a task seems to interact with the appearance of intrinsic activity and motivates further experimental investigation of the functional significance of ongoing activity and associated changes of it.

However, the aforementioned studies are commonly criticised for the possibility that they might confuse “true intrinsic” activity with echoed traces of the previous experiences, i.e. opening the possibility that the observed changes are some kind of memory trace. Naturally, this leads to the question what “true intrinsic” activity

might be – or whether it exists at all. More precisely, this criticism can be applied to any measurement of ongoing activity: whenever resting state activity is measured – be it during wakefulness, during sleep or during anaesthesia – it will always include the so-called task-unrelated mind-wandering (Mason et al., 2007; see Chapter 6 for a more detailed description of mind-wandering) that occurs either with or without awareness of the individual (Smallwood, McSpadden, et al., 2007). Therefore, the very nature of ongoing activity cannot be considered to lack context (Sadaghiani, Hesselmann, et al., 2010). Accordingly, the only possibility to isolate "pure" ongoing activity would require that it possesses a unique spatial and temporal form. As outlined in 1.2.2.4 this does not appear to be physiologically plausible, i.e. no such qualities have been recognised with confidence. More than that, this debate and the question of whether the disputed dissociation of ongoing activity from "other" brain activity is indeed justified or reasonable, leads to the ultimate question about the function of ongoing activity.

Given the observation that ongoing activity does *not* simply represent unconstrained, spontaneous cognition – either called daydreaming, mind-wandering or stimulus-independent thought (Antrobus, 1968; McGuire et al., 1996; Mason et al., 2007) – it seems to reflect a more fundamental or intrinsic property of the brain's functional organisation. In particular, the observation that spatially coherent, spontaneous BOLD activity is present under anaesthesia and during sleep renders it unlikely that the observed patterns of coherence organised in functional networks are solely the result of unconstrained, conscious cognition or mental activity (Christoff et al., 2009; Raichle, 2011). Recently, it has been shown that patterns of low-frequency

oscillations in the BOLD signal are even modulated by the content – or nature – of free thought during rest (Doucet et al., 2012).

Given the research results described in this chapter, one possible function of ongoing activity is the facilitation of responses to external stimuli. Thus, a balance on the global level – similar to the balance of excitatory and inhibitory inputs at the single neuron level that determine the responsiveness of neurons – might be present on a more global level of brain function: opposing forces could enhance the precision of a wide range of processes (Raichle and Snyder, 2007). Indeed, some of these more global effects that involve balance have been reported. For example, the so-called Sprague effect has been first demonstrated in the visual system of the cat (Sprague, 1966).

A more progressive and more expanded view on the functional role of ongoing activity is the notion of predictive coding (see 1.2.1.2) in the context of the experimental investigation of spontaneous fluctuations in the brain. Combining the mentioned electrophysiological findings and the neuroimaging results outlined thereafter, the proposal of the brain a Bayesian interference machine that generates predictions about the future (Olshausen, 2003; Kersten et al., 2004) appears plausible. Most simplistically, the suggestion entails that the brain is shaped by experiences (i.e. stimuli) to represent a best guess, i.e. prior, about states of the environment and – on a cognitive level which holds for humans and some other species – to make predictions about future states of the environment.[2]

2 The question about the initial set of priors equipped with at birth – or even before that – is another interesting one, but shall not be discussed here.

In sum, I propose that the function of ongoing activity is closely related to cognition and that this relation is present during "rest" and during "active states", i.e. it is characteristic for the brain.

1.4 Conclusions

Although the majority of neuroscientists still focus much of their research activities on evoked responses, a growing community of researchers is investigating or taking into account the role of spontaneous or ongoing brain activity. Across all neuroscientific disciplines ranging from single cell recordings in the cat to clinical studies in humans, the importance of spontaneous neural activity is now appreciated and the results add to a rapidly expanding body of research aimed to understand how the brain instantiates behaviour[3]. Raichle and Snyder (2007) even mention the "requirement to establish a framework upon which the study of intrinsic brain activity is incorporated into the work devoted to evoked activity".

Methodologically, the study of slow fluctuations (i.e. < 1 Hz) in neural membrane polarisation has been shown to be of particular interest. In particular, these frequencies correspond to the ones of spontaneous fluctuations in the BOLD signal and their functional consequences seem to be relevant for the understanding of the well-known variability in task-evoked activity, as well as behavioural performance variability.

[3] Here, "behaviour" refers to any mental expression, i.e. a thought is also considered to be a certain type of behaviour.

1.5 This thesis

My interest within the study of the functional role of spontaneous activity is focused on its role in contextual changes ranging from recent experiences of individuals to trial-by-trial variability in a certain task. Therefore, I studied the nature of ongoing activity from two perspectives: One looking at *changes* in the ongoing activity due to learning, and the other one looking at the *predictive role* of prestimulus activity using complementary methodologies, i.e. EEG and fMRI. Finally, I ventured into the realm of inter-individual differences and mind-wandering to investigate the relationship between ongoing activity, certain behavioural traits and neuronal connectivity.

Chapter 2 Methods – measuring spontaneous activity

2.1 Introduction

As mentioned in Chapter 1, several different methodological approaches and techniques can be used to acquire and analyse ongoing brain activity and some of these have been employed in this thesis. This chapter describes the central issues in recording and analysing the data presented, including fMRI and EEG. Other techniques that have not been used here, like electrophysiological recordings in non-human species, are not described to avoid confusion. First, the difference between group studies, lesion studies and inter-individual difference studies is given. Second, a concise overview of the two techniques, i.e. fMRI and EEG, is given and more specific issues related to measuring spontaneous activity are described. With regard to fMRI, this includes a description of the use of stochastic dynamic causal modelling for resting state data.

2.2 Group versus inter-individual differences versus lesion studies

The most common approach to investigate a neuroscientific question in humans using either EEG or fMRI is to perform a group study. In case of the simplest design this entails that healthy volunteers who participate in the study are either randomly assigned to a group, i.e. in case of a between subject design, or do perform two or more experimental tasks or conditions, i.e. in case of a within subject design. For instance, a group study testing for the benefits of a certain drug might randomly assign half of the sample to a group that first takes the drug and then a placebo, and the other half to the reverse order. The gold standard is the so-called double blind

design, where both participant and the experimenter (who interacts with the participant) are “blind” to the experimental condition, i.e. drug versus placebo. A combined between- and within-subject design is also possible. This allows testing for an effect of different dependant variables between and within participants.

One particular form of a group study involves the study of so-called experts, i.e. one group is then comprised of individuals that can perform the studied task(s) particularly well (due to genetic advantages and/or due to intense training).

In comparison to group studies in healthy individuals, clinical studies investigate a sample of patients that show a certain pattern of disease or malfunctioning. Usually, these are then matched, e.g. for variables like age, gender, education and lifestyle, with healthy control participants.

A complementary experimental approach is the study of inter-individual differences. In contrast to group studies, which are based on the assumption that the taken sample is representative of the underlying population and which therefore treat differences between individuals as a source of “noise” that needs to be averaged out, these differences are the main interest for this approach. In the field of differential psychology this is a – or *the* – standard approach to study topics like personality and different types of intelligence. However, only more recently this approach has gained popularity in neuroimaging studies, often linking brain structure to behaviour or personality traits (Kanai and Rees, 2011). More precisely, inter-individual variability from basic to higher cognitive functions including perception, motor control, memory, aspects of consciousness and the ability to introspect can be predicted from structural MRI studies, using voxel-based morphometry (VBM) (Irle et al., 2010;

Kanai et al., 2010; Schwarzkopf et al., 2011) and diffusion tensor imaging (DTI) (Forstmann et al., 2010), as well as from neural activity measured with fMRI (Wig et al., 2008; Bishop, 2009), EEG (Klimesch, 1999) or positron emission tomography (PET) (Gerretsen et al., 2010). Therefore, it has been proposed that these differences can be used to link human behaviour and cognition to brain anatomy and function (e.g. Kanai and Rees, 2011).

Probably the most debated aspect with regard to group and inter-individual difference studies is the group – or sample size respectively – that is needed for a study, i.e. required to show a certain effect. Even though this question is fundamental for any experiment carried out – not only in neuroscience – there are not many publications about the topic (Lenth, 2001). The so-called effect size is a measure of the strength of the deviation from the null hypothesis and usually refers to the estimate of an unknown true effect size based on the collected data (Friston, 2012). One common way to classify effect sizes is to rank them as small, medium and large (Cohen, 1988). Friston (2012) extended this classification to include trivial effect sizes. Essentially, these refer to statistically highly significant effects that are however grounded in "an uninformed overpowered hypothesis test". Based on an analysis of effect sizes in classical inference – which is most often used to report results in neuroscience – Friston (2012) suggests that the optimal size for a sample is between 16 and 32 participants. This argues against the more recent trend of group studies in functional neuroimaging to increase sample sizes due to both editorial requirements and large cohort studies (e.g. Lohrenz, McCabe, Camerer, & Montague, 2007).

Last, there is the option of a case study which is commonly used in psychology and clinical medicine. Sigmund Freud conducted some of the most famous and detailed ones at the beginning of the 20th century, including Little Hans (Freud, 1909a) and The Rat Man (Freud, 1909b). In human neuroscience, case studies are mostly lesion studies, i.e. an individual shows a certain and very specific brain anomaly that cannot be easily compared to a healthy brain. Probably the most famous case is HM and the study of human memory. HM had both of his medial temporal lobes removed and subsequently suffered from intense amnesia (Penfield and Milner, 1958). The benefits, e.g. potential causal inferences can be drawn, and shortcomings, e.g. experimental control, of the lesion study approach have been outlined elsewhere (e.g. Kosslyn & Intriligator, 1992).

2.3 Functional Magnetic Resonance Imaging (fMRI)

2.3.1 Overview

Magnetic resonance imaging (MRI) is a non-invasive method used to create detailed images of the body including the brain by using a strong magnetic field and radio frequency pulses. Instead of creating images of organs and tissues, functional MRI measures blood flow in the brain to detect areas that are active. The technique detects changes in blood oxygenation and flow, which are a consequence of neural activity.

To understand functional MRI, one needs to know how MRI works and how it uses the magnetic properties of tissue: Everywhere in the brain are hydrogen atoms acting as small magnets and rotating around their own axis. When placed in the strong magnetic field (B_0) present in the MRI scanner they stop rotating randomly and align

with B_0. The trick of MRI is to apply a second magnetic field using radio frequency pulses that cause the hydrogen atoms to wobble around their own axis. This movement creates a changing magnetic field around the atoms that in turn creates an electric current. This current constitutes the measured signal of MRI, arising from the whole head. More precisely, when the radio frequency pulse is switched off again, the hydrogen atoms relax, going back to the state of alignment with B_0. Crucially, this relaxation differs between different tissue types, i.e. hydrogen atoms in fat, white matter, grey matter or cerebrospinal fluid have different relaxation times. There are two different relaxation processes that are measured with MRI. First, so-called T1-weighted images focus on the "righting" of tilted hydrogen atoms during the realignment with B_0. This process is influenced by surrounding non-excited molecules. Second, a T2-weighted image is based on the so-called "dephasing" of the rotating atoms, which refers to the falling out of synchrony between different atoms. Mostly, this dephasing is due to the loss of energy to nearby atoms, but can also be influenced by local field inhomogeneities. If the latter factor is taken into account as well, one speaks about T_2*-weighted images. These are the ones acquired with fMRI. Localisation within the three-dimensional space of the brain is achieved via the use of different gradients alongside the radio frequency pulses. Essentially, these gradients generate a magnetic field with different field strength that is changing gradually along three axes: one gradient selects a single slice of the brain and two additional gradients divide the slice into rows and columns allocating so-called voxels. These are essentially cubes of brain tissue from which the signal is collected.

FMRI measures differences in magnetic properties that are related to neural activity. Now, the link to the beginning of this section can be made: Active brain regions

require more blood that transports oxygen and glucose to the active region. Importantly, oxygenated and deoxygenated haemoglobin have different magnetic properties: Oxygenated haemoglobin is diamagnetic and causes no signal loss. For unknown reasons the supply of oxygenated haemoglobin to activated regions is larger than the actual local consumption. Therefore, the proportion of oxygenated to deoxygenated blood is greater. Due to less deoxygenated haemoglobin, dephasing is reduced and a stronger fMRI – or blood oxygen level dependant (BOLD) – signal is obtained (see 2.3.2). In other words, the BOLD signal is the absolute amount of oxyhaemoglobin; the cerebral blood volume (CBV) has to overrule the hyperoxygenation of the haemoglobin. This makes fMRI an indirect measure of brain activity that uses blood flow as a proxy. It has a low temporal resolution, because the response in blood flow lags behind the increase in neural activity and typically peaks between 6 to 8 s after the neural response. The spatial resolution is determined by the applied gradients defining the voxel size and is usually somewhere between 1 and 3 mm, also depending on the field strength of B_0.

Also, BOLD contrast fMRI is a relative measure, because activations need a comparison to a so-called control condition. For instance, if one wanted to localise the brain activation due to moving dots, one could compare a condition of moving dots to a display of static dots. The choice of the control condition is critical: Ideally it should only differ in the parameter of interest (for instance in this case the movement). Typically, fMRI experiments are event-related, blocked or mixed. An event is defined as a short stimulus that is presented to the participant in the scanner, whereas a block is usually constituted of several stimuli that are either identical or

belong to the same category. The BOLD signal of one block will then be compared to a control block consisting of several control stimuli.

2.3.2 The BOLD response

The shape of the BOLD response is characterised by a small initial dip about 1 to 2 s after stimulus onset, a peak at 6 to 8 s poststimulus presentation and a second small dip thereafter, the so-called poststimulus undershoot (Chen and Pike, 2009). The hemodynamic response differs between regions, participants and stimulus presentation rate, but shows a high within-region and -participant stability (Miezin et al., 2000).

The neural mechanisms that underlie the BOLD response as well as the precise mechanisms of neurovascular coupling are not yet fully understood. Glutamate, the main excitatory neurotransmitter, has been identified as a key player in the coupling between blood flow and neurons. When glutamate binds to postsynaptic receptors it causes a calcium influx which activates nitric oxide synthase which triggers the making of nitric oxide (NO). NO leaves the neuron and dilates smooth muscles located around arterioles. This allows more blood to flow to the activated area (Stefanovic et al., 2007). When binding to receptors on astrocytes, glutamate stimulates the production of arachidonic acid and the release of vasodilatory prostaglandins (Takano et al., 2006); vasodilation triggers increased blood flow to the active brain area.

As mentioned above, the transported blood delivers an excess of oxygen, i.e. more than is eventually consumed by the neurons. Although it is still debated why this might be the case, several potential explanations have been put forward. For

example, Fox et al. (1988) proposed that increases in brain activity are supported by the non-oxidative glucose metabolism (i.e. glycolysis).

Maybe the most studied question with regard to the origin of the BOLD signal is its relation to electrophysiological measures. Only if this relationship is understood, neuroscientists are able to compare neuroimaging studies using fMRI with the results of electrophysiological studies. For instance, there is evidence that the BOLD response is correlated to neuronal spiking activity. Comparing single unit data from recordings in monkeys with human fMRI measurements from the motion-sensitive area V5/MT (hMT+), Rees et al. (2000) showed that an increase in motion coherence led to an increase in BOLD responses (in humans) *and* neural firing rates (in monkeys). However, a subsequent study using simultaneous fMRI and electrophysiological recording in monkey visual cortex suggests that the BOLD contrast mechanism may more reflect the input and intracortical processing of a given brain area, rather than its spiking output; the BOLD signal seems to be slightly more correlated to the local field potential then to spiking activity (Logothetis et al., 2001).

2.3.3 How to collect resting state data

In contrast to the typical collection of fMRI data during which the individual in the scanner performs a certain task that is usually projected on a small screen, resting state data is acquired differently. In fact, ongoing activity encompasses the brain activity that is *not* assignable to any given task. Therefore, during collection of resting state data any input due to sensory stimulation should be minimised. Two main approaches are used by researchers: Participants are either asked to fixate on a

central fixation point throughout the scan or to close their eyes without falling asleep. In both cases, it is helpful to tell the participant to try to relax without thinking about anything in particular. Of course, this is an instruction that can hardly be controlled for. In any case, the second option of the above is preferable, given that the participants can be trusted not to fall asleep. Having said so, there is no consensus as to whether the precise experimental setup has a significant impact (Marx et al., 2004; Horovitz et al., 2008; Bianciardi, Fukunaga, van Gelderen, Horovitz, de Zwart, and Duyn, 2009). For instance, the stability of resting state network patterns through various sleep states has been shown to be quite stable, but weaker during deep sleep (Fukunaga et al., 2006; Horovitz et al., 2009). During data collection, the pulse is usually measured, which helps to monitor the physiological state of the participant. In addition, breathing can be measured using a specific respiration belt. These variables can add noise to the signal and should therefore be controlled for (see 1.2.2.2).

The duration of resting state scans differs and commonly ranges between 5 to 10 min (Cole et al., 2010). The chosen duration depends on the planned analysis. While Van Dijk et al. (2010) suggest that 5 min are near-asymptotic with regard to correlation map stability, it appears questionable whether this generalises to analyses that involve a more advanced analysis of connectivity using higher-dimensional independent component analysis (ICA) (Kiviniemi et al., 2009; Smith et al., 2009a), for example[4]. Typically, standard scanning parameters are used including repetition time, echo time, slice thickness, field of view, and matrix size.

[4] For these analyses the degree of partial temporal correlation between subsystems increases. Thus, the ability to delineate them is reduced and more data helps to do so.

2.3.4 Pre-processing and noise correction

Most of the usual pre-processing steps used for task-related BOLD fMRI data are useful for resting state data (Beckmann et al., 2005; Birn et al., 2006). These include realignment of all scans, coregistration to the structural scan of the individual, normalisation to a template brain, smoothing, and correcting for head movements and physiological noise modelling these as regressors during model estimation. Depending on the used head coil, a bias correction taking inhomogeneities in the magnetic field into account might be useful as well.

In addition, there are a number of differences that are worth mentioning. First, high-pass temporal filtering that is usually applied to task-fMRI data might remove parts of the relevant resting state frequency information. Spontaneous BOLD fluctuations have an increasing power at lower frequencies, because they exhibit a 1/f distribution (Fox and Raichle, 2007). Thus, a more conservative correction is necessary to guarantee keeping the power at low frequencies. High-pass filtering is usually performed at 0.01 Hz (Biswal et al., 1995); frequencies above 0.1 Hz mainly relate to cardiac and respiratory factors (Cordes et al., 2001). Thus, spontaneous data acquired with fMRI is often also low-pass filtered at 0.1 Hz leaving a frequency range between 0.01 and 0.1 Hz. Indeed, this range has been shown to be the one in which functional resting state networks are present (Cordes et al., 2001).

Crucially, a large portion of the measured BOLD signal at rest can indeed be attributed to spontaneous BOLD activity. However, artefacts contributing to the signal can be of instrumental or physiological origin. In addition to “real” spontaneous activity, low-frequency changes in the signal might also be caused by

slow changes in head position, physiological parameters and environmental conditions like slow drifts in gain and resonance frequency (Bianciardi, Fukunaga, van Gelderen, Horovitz, de Zwart, Shmueli, et al., 2009). Their relative impact has been investigated using a high field strength scanner, i.e. 7 Tesla[5] (Bianciardi, Fukunaga, van Gelderen, Horovitz, de Zwart, Shmueli, et al., 2009). As mentioned above, low-frequency drifts due to scanner instability contribute to the measured BOLD signal of spontaneous fluctuations (34.7 %), as well as variations in breathing (8.6 %) and cardiac activity (6.6 %). Thermal noise plays a huge role for individual voxel data, but cancels out during averaging, because it is uncorrelated between different voxels. In total, about half of the signal (50.1 %) can therefore be attributed to underlying neuronal activity of interest. These non-neuronal physiological signals can interfere with the interpretation of resting state fMRI data (Birn et al., 2006)

Fortunately, the noise components related to changes in physiological variables, including respiration and heart rate, can be taken care of using the correction method RETROICOR (Glover et al., 2000). It uses the data monitored during scanning by help of the aforementioned breathing belt and a pulse oximeter. This data can then be modelled using a MATLAB toolbox (Hutton et al., 2011). It calculates a time series representing the change in respiration from the mean respiratory waveform by taking the standard deviation at each time point over a 6-second sliding window and convolves it with the so-called respiratory response function (Birn et al., 2008). The resulting respiration volume per unit time (RVT), and basis sets of sine and cosine Fourier series components extending to the 5th harmonic (i.e. 10 regressors) for the cardiac phase and 3rd harmonic (i.e. 6 regressors) for the respiratory phase are used to

[5] Thus, precise numbers given here should be considered with care; they might differ for lower field strengths of 1.5 T or 3 T more commonly used in standard fMRI settings.

model the physiologic fluctuations. This results in a total of 17 regressors which are subsequently added to the first level analyses of each participant's data.

Another potential problem when using ICA is the observed spatial and spectral overlap between some artefactual components and resting state networks. They might even mix and form parts of the same components in the final decomposition (Birn et al., 2008). However, in most cases this overlap seems to be distant enough in order to delineate regions belonging to networks like the DMN and those affected strongly by respiratory fluctuations. This result holds on the single subject basis and at the group level.

Finally, the last pre-processing step that might be included in the analysis of fMRI resting-state studies is the subtraction of spontaneous BOLD fluctuations observed in the whole brain, the so-called global signal. The underlying assumption concerning global signal removal is that the shared BOLD fluctuations are due to physiological factors. There are two ways of taking this into account, thereby ensuring that the mean BOLD signal across the brain is the same in every scan (Fox et al., 2009). Either, the global signal can be added as a linear regressor or multivariate scaling can be used. The scaling is very alike to the approach that is used in task-related studies, i.e. proportional scaling. It removes inter- and intra-session variance in the global signal (Gavrilescu et al., 2002).

Although, the extraction of functional networks from the data can be improved using global signal removal (Fox et al., 2009), it has been shown that this process can induce spurious negative correlations between resting state networks and might result in false positives, i.e. observed anticorrelations (Fox et al., 2005). Thus, observed

results could turn out to be an artefact caused by global signal removal (Murphy et al., 2009). Nonetheless, the nature of the global signal has been shown to be indeed "global", and not restricted or amplified in systems showing anticorrelations. Its removal can lead to an improved correspondence between resting-state correlations and anatomical connectivity (Fox et al., 2009).

In sum, global signal correction should be applied with caution, especially when anticorrelations are observed subsequently.

2.3.5 Functional connectivity analyses of resting state data

As in any other fMRI study, further analyses after pre-processing are guided by the specific type of question that is investigated. In the following, I will describe four different approaches of how to use resting state data, all of which look at functional connectivity. The first one is the seed-based approach. Studies that investigate the relation between spontaneous BOLD activity and any type of behaviour usually focus on one or several brain areas (Hesselmann, Kell, and Kleinschmidt, 2008; Hesselmann, Kell, Eger, et al., 2008). The main benefit of using a seed-based approach in these cases is to be able to ask a straightforward question about connectivity and to generate a direct answer. Due to its simplicity it has been used widely (Cordes et al., 2000; Hampson et al., 2002; Greicius et al., 2003; Fransson, 2005; Margulies et al., 2007; Vincent et al., 2007). A seed-based analysis is performed by first extracting the time course of an *a priori* seed voxel or region of interest (ROI). Subsequently, this time course is used to determine the temporal correlation between the extracted signal and the time courses of all other voxels in the brain (Fox and Raichle, 2007). In order to do so, the data is either used as a

regressor in a linear correlation analysis or in a general linear model (GLM) analysis. The latter one is used in case the model is augmented with confound regressors of no interest, e.g. physiological and movement regressors. This results in voxel-wise functional connectivity maps of co-variance with the seed region (Cole et al., 2010). An assessment of test-retest variability of the method has indicated that resting state network connectivity relationships are identified with moderate to high reliability (Shehzad et al., 2009). However, the technique has several weaknesses.

First, the influence of structured spatial confounds, e.g. other resting state networks than the one under investigation, as well as structured noise, can influence the results. Although some of these influences can be (partially) removed by choosing the necessary pre-processing steps (see 2.3.4), the presence of residual confounds in the seed-region analysis reference time course can negatively influence the resulting correlation maps, because the estimated networks might include voxels that describe the spatial extent of the artefact. Also, the univariate approach neglects the richness of information available within the statistical relationships between multiple data points (Cole et al., 2010).

Second, the selection of time series based on *a priori* hypotheses restricts potential conclusions drawn from the results. One restriction relates to the anatomical confinement of measuring network connectivity, and therefore, on resulting interpretations of hypotheses on the system level. In principle, there are as many possible networks as there are possible seeds, i.e. the interpretation of one spatial map as a meaningful system does not measure up to the data. On the biological level, the choice of the seed might bias any connectivity findings towards specific, smaller

or overlapping sub-systems, compared to larger, distinct networks (e.g. Buckner et al., 2008).

Third and finally, all of these issues depend on the specific choice of the seed. This includes researcher-specific aspects like the size of the seed and its location, as well as subject-specific issues, like functional localisation and spatial normalisation. Testing the potential biases involved in seed-selection performed during seed-based analyses, Cole et al. (2010) compared a number of different version of the DMN, all derived from different versions of seed-based analyses and one estimated using ICA. Independent of the location of the seed voxel, they found significant overlap in the extent of the networks. However, particularly in prefrontal, occipital lobes and subcortical regions, a large amount of variability in the results – and subsequent interpretation – was observed.

In response to the outlined limitations of seed-bases analyses, several other techniques for the extraction of functional networks have been tested. The aforementioned ICA is commonly used (van de Ven et al., 2004; Beckmann et al., 2005; Kiviniemi et al., 2009) and has been used in combination with the more traditional seed-based analysis as well (Seeley et al., 2007). After it had been used to separate uncorrelated signal wave forms in EEG data for several years, ICA was first utilised for the analysis of task-related fMRI data fifteen years ago (McKeown et al., 1998). Only five years later, if was applied to resting state data (Kiviniemi et al., 2003). The basic approach is to decompose a two-dimensional data matrix into the time courses and associated spatial maps of the respective underlying so-called hidden signal sources. In the neuroimaging community, several different approaches

implemented in different software packages have been developed and are in use (Cole et al., 2010). Importantly, the ICA approach has identified the same networks of spontaneous coherence as found with the seed-based analysis approach, including the somatomotor network observed in the pioneering resting state connectivity experiment (Biswal et al., 1995), sensory systems in visual and auditory cortices, as well as those important in higher-level cognitive processes like the DMN. One of the pitfalls of the approach is the manner in which the ICA decomposition is obtained: It uses an iterative optimisation and the order in which the decomposition is performed influences the results.[6] Consequently, a certain degree of run-to-run variability is introduced; therefore, results generated from different ICA analyses can differ between analysis runs on the same data. The use of more stringent convergence criteria can help to reduce this effect. Also, the investigation of the degree of variability can be estimated using specific software (Himberg et al., 2004). Another problem is related to the dimensionality reduction that is performed during ICA: The selection of the number of components is determined by the person who analyses the data. Even though there are algorithms that determine the supposedly optimal number of independent components for given data (e.g. Zuo et al., 2010), no single, "best" dimensionality or model order for the underlying neurophysiology or multiple distributed systems of the brain does exist (Cole et al., 2010). Lastly, the main benefit of the seed-based approach is a major shortcoming of the ICA approach: Whereas a seed-based analysis assures a result that identifies the brain regions most associated – or functionally connected – with the chosen seed, an ICA analysis might be separated into a number of sub-networks (which depend on the parameters of the

[6] Originally, ICA used fixed spatial components. However, more recent versions make use of more variables thereby decreasing the variability of the obtained results.

analysis). Consequently, one might observe a high number of components that are difficult to interpret (Tohka et al., 2008).

The third method used to investigate time series of resting state data is focused on the frequency information contained in the data. Indeed, frequency-specific characteristics have been studied in parallel to correlation-based methods (Cordes et al., 2000). One of the techniques used in order to study these characteristics is the use of so-called amplitude of low frequency fluctuations indices (Zang et al., 2007) and has been discussed elsewhere (Cole et al., 2010).

The last approach, which I want to mention shortly, is the regional homogeneity method (Zang et al., 2004). The technique is sensitive to the homogeneity within clusters that are identified as showing high functional connectivity with a model time series within a certain cluster. Temporal variability within a cluster is indexed by a homogeneity score. Two advantages of this approach are its insensitivity to potential region-to-region or trial-to-trial variability of the hemodynamic response function (see 2.3.2) and the fact that no assumptions about the spatial independence of identified maps are made. Its two main shortcomings are its local nature, which makes the technique vulnerable to different levels of spatial smoothing, and its insensitivity to shape differences between clusters. The latter restriction makes it difficult to draw any inferences on the distributed nature of resting state networks (Zang et al., 2004; Cole et al., 2010).

2.3.6 DCM – or: going beyond functional connectivity

2.3.6.1 Effective connectivity

A different approach to functional connectivity measurements is the investigation of so-called effective connectivity. Generally speaking, there are three different types of connectivity that are investigated using neuroimaging: structural connectivity, which looks at the presence of axonal connections (using for instance DTI, e.g. Taylor and Bushell, 1985), functional connectivity, which has been described above and defines statistical dependencies between regional time series of multiple brain regions, and effective connectivity. The latter one targets causal (directed) influences between neuronal populations (Friston, 1994). Both, functional and effective connectivity are dynamic, i.e. context-dependant recruitment and gating of connections has been observed over milliseconds. Even structural connectivity changes after development has been completed (Friston, 2011).

Effective connectivity goes beyond descriptive statistical methods and requires a causal model. The model defines what is meant by effective direct/indirect causal influences, i.e. it describes the interactions between the elements of the neural system of interest (Stephan and Friston, 2011). These models are defined by dynamic systems theory, a general mathematical framework (Breakspear, 2004; Stephan, 2004). In this framework, a system is defined by its interacting state variables, which are properties that are changing over time. For instance, on the neuronal level the type and number of ion channels that open in response to a stimulation determine the postsynaptic potential. Basically, these functional relationships are depicted by a

number of differential equations[7] including several parameters (e.g. synaptic strength), which determine the type and strength of the causal influences between the different state variables (Stephan and Friston, 2011). In case of the brain, a highly non-autonomous system[8], inputs to the system need to be taken into account.

In sum, the resulting model based on a general state equation that takes state variables, inputs and parameters into account delivers a causal description of how system dynamics result from system structure (Stephan and Friston, 2011); it provides a general model of effective connectivity in neural systems. The completeness is given by including (a) when and where external inputs enter the system, as well as (b) how the state changes in response to inputs over time as given by the structure.

2.3.6.2 Deterministic dynamic causal modelling[9]

The notion of using dynamic causal modelling for neuroimaging data has been introduced ten years ago (Friston et al., 2003), and is one way of measuring effective connectivity. The main difference and advantage of DCM compared to more traditional approaches to measure effective connectivity including structural equation modelling (SEM) (e.g. McIntosh and Gonzalez-Lima, 1991) and psycho-physiological interactions (PPI) (Friston et al., 1997) is the fact that it does not operate at the level of the measured signal. As described above, the signal measured

[7] Differential equations are used because the described systems are referring to a continuous time. A similar approach using discrete timing and difference equations has been used as well (e.g. Harrison et al., 2003).

[8] A system is considered non-autonomous if it exchanges energy, matter or information with its environment.

[9] Before stochastic DCM had been developed, there was only one DCM approach. In order to differentiate between different DCM approaches, the "classical" DCM got renamed and is now commonly referred to as "deterministic DCM".

in humans during neuroimaging is indirect, but the causal architecture that is targeted is at the neuronal level. Thus, two models are required: one that links neuronal activity to the measured haemodynamic signal, and one that models neuronal population dynamics (Stephan et al., 2004). These models optimally fit the parameters to maximise similarity between the predicted (modelled) time series and the measured (observed) time series. To my knowledge, DCM is the only approach that incorporates an obligatory combination of models of neural dynamics (i.e. the state evolution model) and biophysical forward models (i.e. the observation model) (Stephan and Friston, 2011).

As mentioned in 2.3.6.1 and in order to calculate the state evolution model, DCM makes use of a neural state equation that takes the system's state (i.e. the neural state variables), the inputs and some parameters that define the functional architecture and interactions between brain regions at the neuronal level into account. For each brain region, the change of a neural state vector in time is modelled as a single state variable by use of a bilinear differential equation. The resulting values represent a summary index of neural population dynamics, i.e. mean regional activity. The neural dynamics are driven by experimental inputs specific to the respective study. They can affect the model in two different ways. On the one hand, they can influence specific regions directly (e.g. evoked neural responses in early visual cortex due to a visual stimulus) or they can modify the connections between different regions, i.e. the coupling or connection strength between brain areas, e.g. after learning has occurred or via attentional mechanisms. Usually, the latter ones are of main interest for a related scientific question.

In order to model haemodynamics (i.e. the observation model), DCM uses the so-called “Balloon model”, which has been validated (Buxton et al., 1998; Friston et al., 2000; Stephan et al., 2007). Put shortly, it combines several differential equations that use a set of parameters to describe how fluctuations in neural activity cause changes in blood flow and volume, as well as the vasodilatory signal and the content of deoxyhaemoglobin. The resulting prediction of the BOLD signal is a non-linear function of blood volume and the content of deoxyhaemoglobin (Stephan et al., 2007).

To summarise, both neural and haemodynamic parameters are estimated on the basis of the acquired BOLD data. In order to do so, a Bayesian approach is used (Friston, 2002). Different priors are used for the haemodynamic parameters and the coupling parameters.

After estimation, the parameters are used to draw inferences about the neural mechanisms underlying the task at hand, i.e. whether and if so, where and how the experimental design resulted in change(s) of connection strength between different regions of the tested network(s). On the group level, comparisons can be performed in different ways. Similar to a random effects analysis, the parameter estimates of interest can be compared in a classical second-level analysis between individuals, i.e. a t-test of the parameter(s) of interest (Smith et al., 2006). Corresponding to a fixed effects analysis – and thereby restricting the scope of potential inferences to the sample investigated – Bayesian statistics can be used (Garrido et al., 2007).

2.3.6.3 Stochastic DCM

How could DCM be used for resting state data? Essentially, such data does not have any input by definition. Also, the endogenous fluctuations in neuronal and vascular responses need to be modelled somehow. Stochastic DCM takes care of both of these challenges and can be applied to task-driven and resting-state data. It uses variational Bayesian generalised filtering, which has been described elsewhere (Friston et al., 2010; Li et al., 2011). The challenge of stochastic DCMs is trifold: In contrast to deterministic DCMs that need to estimate two sets of variables, i.e. the parameters and the hyperparameters, stochastic DCMs include a third set of unknown variables, i.e. the hidden states. Recently, Li et al. (2011) tested the validity of stochastic DCM for the use on fMRI data. They compared results of data with high and low levels of physiological noise and proposed that stochastic DCM has face and construct validity. They conclude that the ability to model spontaneous fluctuations in hidden neuronal states provides a new perspective on the regionally specific generation of fMRI signals.

Regions – or nodes – for the chosen stochastic DCM are extracted in a similar way as for a deterministic DCM. The selection of the chosen coordinates can be either based on prior hypotheses using anatomical atlases or on functional connectivity analyses. The latter one might include distinct clusters from seed-correlation analyses or the results of an ICA (see 2.3.5).

The main challenges of stochastic DCM are related to the selection of the nodes: First, strong hypotheses about which ROIs are included in the analysis, which topography might be tested and which aspects (e.g. existence of connections versus

the modulation of connections) are tested, are required. Second, the definition of ROIs in different individuals depends on multiple decisions and is therefore error-prone and less reproducible.

2.3.6.4 Model selection

Model selection is integral to both, deterministic and stochastic DCM. Put simply, model comparison – which is involved in almost every form of scientific reporting – tries to answer the question which of several alternative models represents some observed data best. Two factors should be considered when comparing models: fit and complexity. While the first aspect is obvious, the issue of complexity is less commonly thought about. However, it is as crucial and can be best explained in the context of so-called over-fitting, which occurs when a model gets more and more complex, usually fitting noise present in the data. The difficult decision during model fitting and subsequent model selection is to decide which fit is *just* sufficient to explain the data in an optimal way without over-fitting the data.

Model selection between different models within a DCM analysis is based on Bayesian model selection (BMS), which compares the Bayes factor of two models with each other using an estimate of the log evidence for each model (Stephan and Friston, 2011). It is a crucial part during the DCM analysis and is commonly applied (e.g. Grol et al., 2007; Kumar et al., 2007; Acs and Greenlee, 2008) before inferences about particular parameters are drawn.

Another approach of using model selection compares two or more subsets of model space (Stephan et al., 2009). By choosing the subsets such that they reflect the components of the model structure one is interested in, this approach goes beyond

the comparison of different models. For instance, one can test whether a certain connection is included in the optimal model or whether a specific connection is modulated by the experimental condition. The main advantage of this approach is that it allows the comparison between large sets of models simultaneously.

2.4 Electroencephalography (EEG)

The first human EEG was recorded in 1924 (Berger, 1929). EEG measures electrical activity along the scalp or more precisely, it records neuronal voltage fluctuations due to ionic current flows. In comparison to fMRI it has a high temporal resolution (commonly used sampling rates are between 250 and 2000 Hz), because it measures a direct neural signal. However, it suffers from a poor spatial resolution, because the electrodes are attached to the scull and the signal measured at each electrode is a weighted sum of neuronal activity from different sources. Also, some currents produce potentials that cancel each other out.[10] This challenge of data interpretation is called the inverse problem and describes the mathematically impossible task to reconstruct a unique intracranial current source for any measured EEG signal. Related to this, neuronal activity from subcortical areas can only be measured poorly.

EEG signals are analysed in terms of rhythmic fluctuations, which are divided into frequency bands of brain oscillations: the delta (2 – 5 Hz), theta (5 – 8 Hz), alpha (8 – 14 Hz), beta (14 – 25 Hz), and gamma (25 – 100 Hz) band[11]. The measured fluctuations are attributed to local field potentials between excitatory and inhibitory

[10] The EEG signal really only represents a specific type of neuronal activation, i.e. the activity in aligned (elongated) dendrites that are all activated in the same way. In other words, the technique depends on a very specific coherent neuronal activation.

[11] The division between frequency bands (still) differs between different scientific accounts, because they are not standardised.

states of neural populations. Oscillations are the summated excitatory post-synaptic potentials (EPSPs) of several thousands of neurons (Lopes da Silva et al., 1980). The raw EEG signal contains various different brain oscillations which can be analysed using spectral analysis (e.g. wavelet analysis, Başar et al., 2001). In order to characterise an oscillation unambiguously, two more parameters next to the frequency are needed, i.e. amplitude and phase.

The raw EEG power can distinguish between gross changes in state like alertness. For example, alpha oscillations are associated with a relaxed, but awake state of the participant while beta oscillations are associated with a more alert and active state. Even though this analysis in terms of gross changes in state gives certain ideas about the mental state of an individual, it does not allow addressing specific, more fine-grained changes in mental activity. Therefore, its use in the study of moment-by-moment human cognitive activity is very limited.

2.4.1 Event-related potentials (ERPs)

Event-related potentials (ERPs) allow the investigation of short events such as used in typical perceptual, motor or cognitive experiments in human neuroscience. The basic idea is to time-lock the recording of the EEG to the stimulus onset and to average the response over multiple trials. This allows filtering out the brain activity unrelated to the presentation of the stimulus; averaging over a large number of trials the random activity in the signal approaches zero as the number of trials increases. The waves that survive this averaging process reflect deviations from a prestimulus baseline and are referred to as ERP components. The amplitude of an ERP component is relatively small, ranging from less than 1 to 10 microvolts, in

comparison to the raw EEG ranging between 10 to 100 microvolts. The waves possess several positive and negative voltage deflections. These are related to a set of underlying ERP components (Luck and Kappenman, 2012). Most of these components are identified by their polarity (P for positive and N for negative) and a number indicating the latency in milliseconds or the position in the waveform respectively. For instance, the language-related N400 has a negative polarity and occurs around 400 ms after stimulus onset. However, the latencies often encompass large intervals, e.g. the P300 may peak between 250 and 700 ms (Luck, 2005).

Traditionally, ERPs are supposed to be superimposed on ongoing EEG random activity (i.e. noise) consisting of an amplitude and a phase distribution unrelated to the processing of the stimulus. However, this view has been challenged more recently. First, ERPs might represent phase resetting of the ongoing EEG triggered by the stimulus which leads to transient time- and phase-locking of frequency-specific oscillations to the stimulus onset of every trial (Makeig et al., 2004). These phase-synchronised oscillations survive the averaging over trials and become visible as waves in the average ERP. Second, time-frequency analyses of single trial EEG epochs have revealed that EEG does not merely consist of random background noise. Instead, event-related changes in the magnitude and phase of EEG oscillations at specific frequencies have been observed (Makeig et al., 2004).

2.4.2 Time-frequency analyses (TFAs)

Generally speaking, time frequency analysis (TFA) includes those signal processing techniques that study a signal in two domains: time and frequency, i.e. the spectral and the temporal domain. In order to visualise the results of a TFA, a time-frequency

plot or spectrogram is used. This shows the temporal evolution of the spectrum in a 2D manner, where the color-coding provides information about the power of a given frequency band at a given time.

TFA complements the EEG signal analysis in several ways. First, it takes into account that neurons are oscillating. Second, it analyses the signal including the trial-to-trial jitter, which might contain important information. Third, it allows the analysis of longer time periods including prestimulus activity and spontaneous signals. With regard to this thesis, the third aspect is of special importance.

Most commonly, TFA are calculated using Fourier analysis and wavelet analyses. The Fourier transform (FT) converts a signal from the time domain to the frequency domain (Mitra and Pesaran, 1999). A time-frequency representation of power is calculated by use of a sliding time window. This either has a fixed length independent of frequency, or a decreasing length with increased frequency. In both cases the power is calculated for each time window. Prior to that one or more tapers are applied to the data to reduce spectral leakage and to control the frequency smoothing. Wavelet transforms co-localise in both domains, i.e. time and frequency, and can be used for non-stationary signals like the one acquired with EEG (Mitra and Pesaran, 1999; Başar et al., 2001).

Chapter 3 Early visual learning induces long-lasting connectivity changes during rest in the human brain

3.1 Introduction

As outlined in Chapter 1.1, until recently functional MRI (fMRI) studies have focused on how brain activity changes with task performance or sensory stimulation. However, even at rest – in the absence of a task or stimulation – fMRI signals show spontaneous fluctuations that exhibit spatiotemporal correlations in networks of functionally connected areas (Biswal et al., 1995; Fox and Raichle, 2007; Raichle, 2010). These networks continue to covary during sleep (Fukunaga et al., 2006) and under anaesthesia (Vincent et al., 2007). They show high consistency and reproducibility across subjects and sessions over the short-term and long-term, using different variations of independent component analysis (ICA) (Damoiseaux et al., 2006) and group ICA (Zuo et al., 2010). Their reproducibility in healthy young individuals compares to that of activations elicited by motor paradigms (Meindl et al., 2010). Furthermore, there is a close correspondence between the activation networks – of almost 30,000 human participants of fMRI studies – with resting state networks (Smith et al., 2009b). The interplay between spontaneous and evoked activity has been of particular interest. For example, in the visual cortex, spontaneous fluctuations determine the variability in cortical responses and perception associated with presentation of a simple visual stimulus (Schölvinck et al., 2012).

The effect of spontaneous fluctuations on evoked responses associated with perception raises the complementary question of whether systematic changes in

evoked responses, for example present during learning, might subsequently alter spontaneous fluctuation. The mechanism I have in mind here is that experience dependent (associative) plasticity may change synaptic connections and ensuing neuronal activity in the local circuits affected. As the implicit short term and immediate long-term potentiation is consolidated the associated changes in spontaneous neuronal activity should persist and be measurable in terms of changes in effective connectivity. A growing number of studies have investigated this adaptive modulation of resting state networks (see 1.3). Changes in spontaneous fluctuations have been shown after visuo-motor learning (Albert, Robertson, and Miall, 2009), episodic memory tasks (Tambini et al., 2010), and language tasks (Hasson et al., 2009).

Visual learning is one way in which systematic changes in cortical responses and perception can be induced. Intensive training on a simple shape identification task over several days can change resting state functional connectivity between visual and fronto-parietal cortices (Lewis et al., 2009). This indicates that visual learning can have lasting effects on spontaneous brain activity through experience dependent plasticity. But such effects occur only after several days of training. The early phase of visual learning occurs much more rapidly—and is often ignored in typical visual learning experiments. However, learning entails a rapid consolidation process that starts within a single training session (Seitz and Watanabe, 2005) and that occurs in any experiment, independent of modality. The specific changes in spontaneous activity in task-responsive brain areas in response to this early learning (that occurs in any experiment, independent of modality) perhaps more typical of real-world environments (Brovelli et al., 2008; Shtyrov, 2012) remain unknown. With regard to

visual learning, both sensory and non-sensory areas (Goldstone, 1998; Seitz and Watanabe, 2005; Adab and Vogels, 2011; Shibata et al., 2011), appear to be involved. Outside the sensory cortex, single-neuron and functional MRI studies have implicated the lateral intraparietal area (Law and Gold, 2008), lateral parietal cortex (Kahnt et al., 2011), subcortical structures like the hippocampus (Lee et al., 2005; Graham et al., 2006) and the caudate nucleus (Ding and Gold, 2010). Recently, sub-areas of the medial temporal lobe (MTL) including the parahippocampal cortex and subiculum have been implicated in rapid and incidental statistical learning in a visual paradigm (Schapiro et al., 2012). While MTL regions and, importantly, the hippocampus – including its connections to the striatum – have been traditionally linked to memory processes such as memory consolidation (Battaglia et al., 2011), their role in perceptual learning has only been examined more recently (Buckley, 2005).

Memory consolidation refers to the processes underlying the stabilisation of memory traces acquired during initial encoding (Dudai, 2012); where the importance of sleep for consolidation is well-established (Wang et al., 2011). Previous studies of changes in resting state activity in response to recent experiences (Albert, Robertson, and Miall, 2009; Lewis et al., 2009; Tambini et al., 2010) have not examined long-term changes in spontaneous fluctuations in the resting state. This is probably due to the fact that this requires a more extensive study design. However, I was particularly interested in these potential long-term changes as markers of experience dependent plasticity induced by the early learning phase.

Therefore, I used a paradigm with only one relatively short learning session that promoted rapid perceptual learning. I hypothesized that rapid perceptual learning would be accompanied by changes in spontaneous activity in brain structures whose responses changed during learning. Furthermore, I predicted that resting state changes would persist following consolidation. I tested this hypothesis by acquiring resting state time-series using functional MRI before and after a standard perceptual learning experiment. During the experiment, participants learned to discriminate a motion stimulus. I measured brain responses during task learning to identify regions whose responses were correlated with the learning in each individual. Crucially, I also acquired independent measures of resting state brain activity before and immediately after learning. The following day, I repeated the paradigm without the learning. I used stochastic DCM to evaluate resting state effective connectivity (Li et al., 2011) between regions identified in the learning session. Specifically, I tested for learning dependent changes in effective connectivity (during rest) immediately after the learning session and after consolidation of these putative changes on the following day.

3.2 Materials and methods

3.2.1 Participants

16 right-handed healthy volunteers (7 female, 19 – 33 years of age, mean age 25.4 years) with normal or corrected to normal vision gave written informed consent to participate in the study consisting of two scanning sessions at two consecutive days. 11 of the 16 participants learned the motion coherence task and were included in the data analysis (3 performed at ceiling level and were excluded because I did not

expect to see any neural changes in the absence of behavioural improvement; 2 were not able to learn the task as disclosed by their persistently low performance). The study was approved by the local ethics committee.

3.2.2 Stimuli and task design

A random dot motion coherence stimulus was used. The level of dot motion coherence was set to 20 %, which is close to the perceptual threshold and has been successfully used for naïve participants previously (Vaina et al., 1998). Further stimulus parameters were chosen according to the results of a behavioural piloting study of 15 participants. All participants performed 25 task and 25 control blocks, each consisting of 16 trials. A presentation time of 0.3 s was used in 7 subjects and 0.4 s in the remainder. The longer presentation time resulted in ceiling performance for 6 of the 8 participants. Therefore, I chose a 0.3 s presentation time for the scanning paradigm. The following stimulus parameters were used: dot speed: 10 °/s, dot life time: 6 frames, response time: 1.5 s, number of dots: 200. White dots were presented at maximum contrast in a central circular aperture covering a 3.14 ° visual angle on a black background. Participants were asked to focus on a white fixation square at the centre of the screen throughout the experiment and no feedback was given.

During trials of the motion task, 80 % of the dots were moving in random directions across the screen, while 20 % of the dots were moving coherently to the left or right. The coherent direction was chosen randomly. Participants used their right hand and a keypad to report the direction of motion; i.e. left or right after the stimulus had disappeared. During control trials the dots were static and a little arrow, pointing

either to the left or right, replaced the central fixation square. In these trials participants reported the direction of the arrow.

In total, each participant completed 800 trials—400 trials of the motion learning task and 400 trials of the control task, divided into 25 blocks of 16 trials each. The 25 blocks were spread over 5 runs, i.e. the scanner was restarted after 5 blocks – allowing participants to rest between runs. Each block of the motion task was followed by a block of the control task or vice versa.

3.2.3 Experimental procedure

To address potential changes in resting state connectivity due to learning and consolidation, participants were scanned on two consecutive days and brain signals were measured in four resting state runs: one before task performance, one directly after task performance, and two at the second day. These were repeated at the same times as the rest runs on the first day. Participants underwent standard retinotopic mapping and a V5/MT localizer in the scanner between the two rest runs of the second day (see Figure 3-1). Before entering the scanner on day one, participants were familiarized with the task, but did not pre-train (to ensure perceptual learning during scanning). Task instructions emphasized that accuracy was more important than speed when responding. Both scanning sessions lasted about 90 min (see Figure 3-1 for details), and were separated by 24 h for each participant. During resting state runs participants were asked to close their eyes, relax, and to not fall asleep. The order of motion and control conditions in the learning task was counterbalanced over subjects.

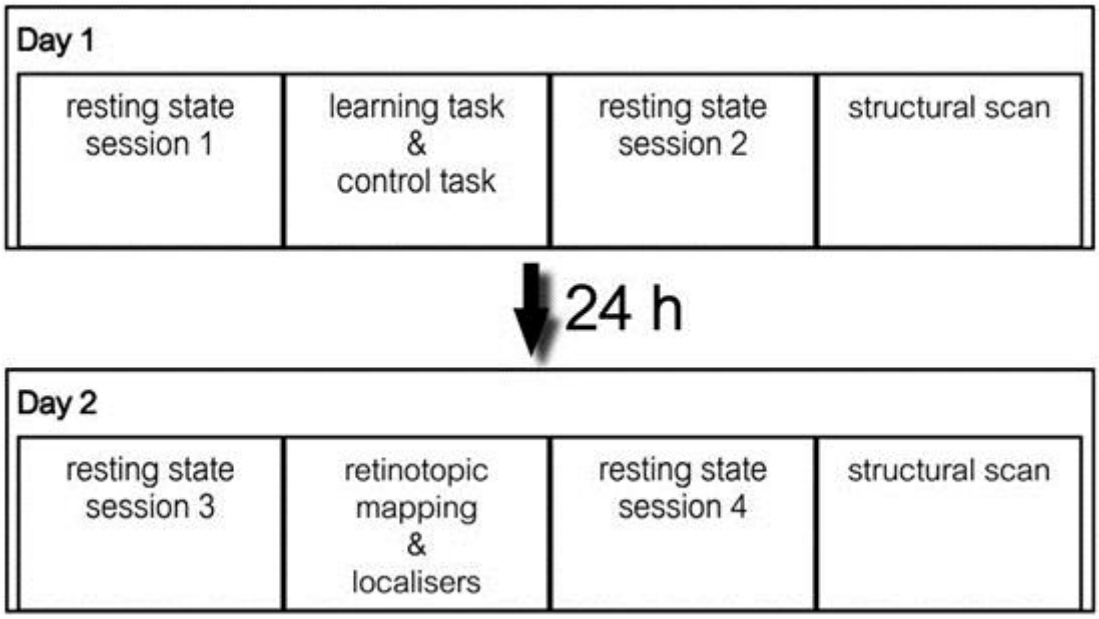

Figure 3-1 Experimental paradigm. Participants were scanned on two consecutive days for about 90 min each day. Two resting state runs were acquired each day, preceding and following the learning task or a retinotopic mapping respectively. A structural scan was acquired on both days.

3.2.4 Behavioural analysis

Behavioural data were analysed using inverse efficiency (IE)—a simple measure that combines reaction time and accuracy; where IE = mean reaction time / accuracy (Graham et al., 2006). Single trial reaction times that deviated from the mean of the respective block by more than three standard deviations were excluded. IE was calculated for each block (n = 25) and raw values were fitted to an exponential function of the form $y = ae^{-bx}$ where a represents the amplitude and b the learning rate. The ensuing estimates of inverse efficiency were used as a parametric modulator of the stimulus regressors in the first level (within-subject) analysis of the functional data acquired during the learning task (see below). These regressors modelled learning related adaptation of BOLD responses.

3.2.5 fMRI data acquisition

A 3 T Trio MRI Scanner (Siemens Medical Systems, Erlangen, Germany) with a 32 channel head coil was used to acquire functional data with a standard echo planar imaging (EPI) sequence (matrix size 64 × 64; field of view 192 × 192 mm; in plane resolution 3 × 3 mm; 32 slices in ascending acquisition order; echo time 30 ms; acquisition time per slice 68 ms; TR 2.176 s). Each run of the learning task comprised 246 volumes, and each resting state acquisition comprised 276 volumes. On both scanning days, B0 field maps were acquired to correct for geometric distortions in the EPI images. Also a structural T1-weighted scan was acquired on both days (matrix size 256 × 240; field of view 256 × 240 mm; in-plane resolution 1 mm × 1 mm; 176 sagittal slices of thickness 1 mm; echo time 2.48 ms; acquisition time per slice 7.92 ms). During scanning, respiration volume and cardiac pulse were measured using a breathing belt placed around the participants' waist and an MRI compatible pulse oximeter attached to one of the fingers. These data, together with scanner slice synchronization pulses, were sampled using Spike2 (Cambridge Electronic Design Limited, Cambridge, UK) and used for physiological noise correction.

3.2.6 fMRI data analysis

Functional data were analysed using SPM8 (http://www.fil.ion.ucl.ac.uk/spm/software/spm8/) and DCM12 was used for dynamic causal modelling of effective connectivity. To allow for T1 equilibration, the first five images of each run were discarded. Pre-processing of the data involved mean bias correction, realignment of each volume to the first volume of each run,

coregistration of the functional data to the structural data of each day, coregistration of the structural scan (and functional volumes) of the first day to that of the second day, normalization to the MNI template brain and smoothing by an 8 mm Gaussian kernel. The task data were filtered with a standard 128-s cut-off and the resting state data were filtered with a 256-s cut-off, high-pass filter to remove low-frequency drifts – including differences between runs, while preserving as many of the spontaneous fMRI fluctuations as possible (see 2.3.4; Birn, 2007). Physiological data (respiration and heart beat) were modelled using an in-house developed MATLAB toolbox (Hutton et al., 2011) based on RETROICOR (Glover et al., 2000); see 2.3.4. This resulted in a total of 17 regressors. The resulting regressors were included as confounds in the first level analysis for each participant. Movement parameters were also included as confounds. No global signal regression was performed.

3.2.6.1 Perceptual learning session

Regressors modelling the stimuli were formed by convolving boxcar functions encoding each condition with a canonical hemodynamic response function – where stimulus functions modelling learning blocks were parametrically modulated by the fitted values of inverse efficiency (IE). These stimulus functions model perceptual learning related changes in responses evoked during the learning task. Contrasts of first level parameter estimates were used to perform a random effects analysis over participants in the usual way. This involved estimating (contrasts of) parameters encoding the effects of interest using a standard linear convolution model at the first (within-subject) level (over all five task runs) and then passing the resulting contrast images to one sample t-tests at the second (between-subject) level. The resulting statistical parametric maps (SPMs) were used to test for differences between the

learning and the control task, the learning task and the fixation baseline, and the effects of learning (i.e. testing for a parametric modulation of the learning task effects). The anatomy toolbox (Eickhoff et al., 2005) was used to anatomically designate activated areas.

3.2.6.2 Psychophysiological interaction analysis

The peak activation – elicited by the effect of perceptual learning – was used as a region of interest (ROI) for the analysis of the resting state data. Time series of this ROI were extracted for all four resting state runs and included as regressors in a first level general linear convolution model, together with the nuisance regressors. Again, resulting contrast images were passed to one sample t-tests at the second (between-subject) level and the resulting SPMs were used to test for changes in the coupling with the region defined during the learning task. More precisely, the four rest runs constituted two main effects, i.e. the main effect of day (rest 1 and 2 vs. rest 3 and 4) and the main effect of time (rest 1 and 3 vs. rest 2 and 4). The interaction of the two effects, i.e. day × time, was used to test for changes in the coupling between the learning related ROI and any other brain region (regression slope of regional activity on the activity of the ROI). Participant-specific peak coordinates of the learning related region were used. The peaks ($p < 0.05$, uncorrected) were within 16 mm of the second-level (between subject) peak and within the specific anatomical region, as defined by the SPM Anatomy toolbox (Eickhoff et al., 2005). Together with the learning related region, the region showing the most significant psychophysiological interaction (over subjects) was used for the subsequent DCM analysis of changes in their effective connectivity.

3.2.6.3 Dynamic causal modelling

DCM models neuronal dynamics in terms of directed and reciprocal influences among brain regions. Stochastic DCM allows one to model spontaneous or endogenous (non-controlled) activity. It does not require any input usually associated with experimental manipulation. Two participant-specific ROIs defined by the learning task and the psychophysiological interaction analysis were used as the nodes for 10 different models of changes in extrinsic connectivity. Regional activity in each ROI was summarised with its principal eigenvariate, adjusted for nuisance variables, based on voxels within 8 mm of subject-specific peaks. All four runs were concatenated into a single time series and parametric modulators were used to model learning-related changes in effective connectivity, plus potential consolidation of these changes.

More precisely, run-specific differences – in terms of the (bilinear) modulation of the average connectivity over all four rest runs – were modelled with three different parametric modulators. First, I modelled non-specific adaptation (i.e. the effect of "run") due to time in the scanner by weighting the four different rest runs accordingly by [0 1 0 1]. Second, I added the effects of visual learning – following the learning phase – using the following weights [0 1 0 0]. Finally, a consolidation model comprised adaptation effects, i.e. [0 1 0 1], and learning effects that persisted during the second day with the following weights [0 2 1 1]. Crucially, the learning and consolidation models have two bilinear coupling parameters per connection that control the relative expression of adaptation and learning (or consolidation) respectively. I applied the models of coupling changes, – including a null model with no changes in coupling – to different permutations of connections: forward

connections from one region to another, backward connections from one region to another, and bilateral connections, involving both forward and backward connections. This resulted in models with the same extrinsic reciprocal connections between two nodes, but different modulations of those connections. All models were fitted to the concatenated time series of the rest runs using generalized (Bayesian) filtering (Li et al., 2011). To evaluate the relative evidence for each of the 10 models, I compared the (variational free energy approximation to) log evidence. I used Bayesian Model Selection to select the model with the greatest evidence given the data. More precisely, I used relative log evidences, i.e. the model with the least evidence was subtracted from each model. This fixed effects model comparison was used because I assumed that the same model accounted for the data generated by every participant. A difference of three between log evidences – which corresponds to a relative evidence (Bayes factor) of about 20:1 – was used as the criterion for model selection.

For quantitative interpretation, the changes in effective connectivity under the winning DCM were computed by multiplying the appropriate bilinear parameters with the run-specific weights as specified above. Thus, for each participant each connection between the two regions included in the model was described by four values, reporting the connection strength in each resting run, relative to the first.

Non-specific adaptation between the first and second scanning day was not modelled, because I assumed that resting state connectivity would not show cumulative changes over successive days when learning had only occurred on the first day. Furthermore, I emphasise that the consolidation model did not simply

represent a non-specific change in effective connectivity on the second day: it had to change in proportion to the learning-dependent changes on the first day.

3.3 Results

3.3.1 Participants showed early rapid learning of the motion task

Participants completed 400 trials of the motion task and 400 trials of the control task. Performance was measured using inverse efficiency (IE). The IE values of each block were fitted with an exponential function. See 3.2.4 for details. See Figure 3-2 for an overview of the learning. The fitted IE values entered the analysis of the functional neuroimaging data as a parametric modulation of the stimulus regressors in the first level (within-subject) analysis of the learning run.

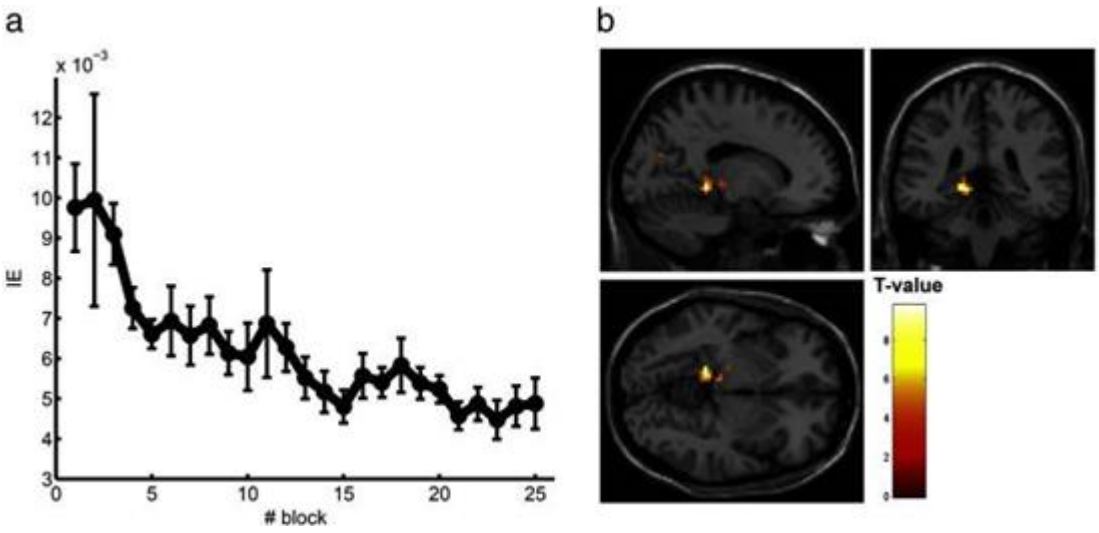

Figure 3-2 Behavioural learning and hippocampal activation. a) Participants learned the motion task. Inverse efficiency (IE) is plotted for every block of the task (n = 25). Data are averaged over all participants (n = 11) who learned the task successfully. Error bars show SEM. b) Learning activated the hippocampus. The fitted inverse efficiency values of the learning task were used for the plotted contrast. Statistics were significant at $p < 0.05$, FWE corrected. Images show activation at $p < 0.001$ (uncorrected).

All participants who learned the task performed (as expected) at ceiling on the control task throughout the 25 blocks (mean of all participants over all blocks: 99 % correct, range between participants: 97 % to 100 % correct).

3.3.2 Motion task activated visual, frontal and parietal areas

After pre-processing, I first identified regions showing activity specific to the motion task by contrasting the blocks when participants performed the motion task with the fixation baseline. I found a bilateral network of visual areas, including V5/MT, as well as inferior parietal and orbitofrontal cortex (all $p < 0.05$, FWE corrected). See Table 3-1 for an overview.

Next, I examined activations associated with the motion task compared to the static control task and found these in the inferior parietal cortex and the right insula cortex (all $p < 0.05$, FWE corrected), as well as in the visual cortex extending into V5/MT and medial temporal regions, and in the medial frontal cortex (all $p < 0.001$, uncorrected). See Table 3-2 for an overview.

Table 3-1 Main effect of the motion learning task compared to baseline. Voxel-level statistics are reported at $p < 0.05$, FWE corrected. BA = Brodmann area, L = left hemisphere, R = right hemisphere.

	MNI coordinates				
	x	y	y	t-value	P-value
BA 18 R	− 24	− 94	13	15.43	0.001
BA18 L	15	− 91	− 2	14.48	0.001
Fusiform gyrus R	36	29	− 2	11.35	0.011
Inferior parietal cortex L	− 30	− 46	49	10.58	0.021

Inferior occipital cortex L/MT	− 45	− 67	-2	10.17	0.030
BA 17/cuneus R	12	− 94	13	10.10	0.032
Inferior orbital frontal cortex L	− 42	20	− 2	9.96	0.036
Medial occipital cortex L	− 42	− 76	7	9.85	0.040
Inferior parietal cortex L	− 27	− 43	40	9.78	0.043

Table 3-2 Main effect of the motion learning task compared to the static control task. Voxel-level statistics are reported at [a]$p < 0.05$, FWE corrected or [b]$p < 0.0001$, uncorrected, BA = Brodmann area, L = left hemisphere, R = right hemisphere.

	MNI coordinates				
	x	y	y	t-value	P-value
Inferior parietal cortex R	36	− 37	34	11.81	0.007[a]
Insula R	33	29	1	10.12	0.031[a]
BA 18/19 L	− 24	− 79	13	9.31	P < 0.0001[b]
Inferior parietal cortex L	− 36	− 37	40	9.31	P < 0.0001[b]
Medial cingulate cortex R	9	17	47	8.14	P < 0.0001[b]
Insula L	− 36	20	− 5	7.61	P < 0.0001[b]
Precentral sulcus R	27	− 7	55	6.91	P < 0.0001[b]
Medial frontal cortex R	45	35	31	6.78	P < 0.0001[b]
Inferior frontal gyrus L	− 57	14	28	5.86	P < 0.0001[b]
Caudate nucleus R	15	− 4	19	5.13	P < 0.0001[b]

3.3.3 Early learning-related modulation of hippocampal activity during task performance

Using the IE-based parametric regressor, which modeled participant-specific learning on the motion task, I tested for regions whose responses adapted with performance. This analysis identified the left hippocampus (left subiculum, MNI coordinates (x = − 15, y = − 37, z = − 5), t = 9.77, p = 0.04, FWE corrected) (see Figure 3-2b). The anatomy toolbox assigned the activation to the left subiculum with a 100% probability. None of the motion-activated areas given in Table 3-1 and Table 3-2 showed any learning related changes ($p < 0.001$, uncorrected).

3.3.4 Learning-related changes in connectivity during rest

Having identified the hippocampus as the key region whose activity changed significantly with perceptual learning (as indexed by participant-specific changes in performance) I next explored how the resting state connectivity of this region changed after learning. I first identified candidate regions whose connectivity with the hippocampus changed between resting state runs using a psychophysiological interaction analysis (Friston et al., 1997). These regions were then used in a dynamic causal model to examine changes in effective connectivity with the hippocampus. Using the independently acquired resting state data, I extracted the time series of the participant-specific hippocampal peak voxels for all four resting state runs and tested for changes in the coupling of the hippocampal region of interest with learning using PPIs.

To test for these changes, I treated the resting state runs as a 2 × 2 factorial design. Testing for the interaction between the two main effects of "run" (i.e. run one and

three vs. run two and four) and "day" (i.e. run one and two vs. run three and four) I found that bilateral striatal loci showed changes in coupling with the hippocampus between runs that were significantly greater on the first compared to the second day (MNI coordinates (x = − 21, y = 8, z = 2), t = 3.59, p = 0.002, uncorrected; MNI coordinates (x = 21, y = 14, z = 4), t = 4.94, p < 0.001, uncorrected). No other regions showed a run by day interaction (p < 0.001, uncorrected) and I used the striatal region for the dynamic causal modelling.

3.3.5 Dynamic causal modelling

My subsequent tests for learning-related changes in effective connectivity (i.e. plasticity), and potential consolidation of these changes, were based on Bayesian model comparison using stochastic DCM (Li et al., 2011). My models differed in terms of when and where changes in connectivity were expressed, i.e. specifically characterising the forward and backward connections between the left hippocampal and striatal regions identified by the conventional SPM and PPI analyses. My hypotheses were not about the existence of connections, but whether there were changes in specific connections between these areas across the different rest runs. Therefore, I considered four types of models: first a null model without any changes in connectivity (null). Second, I considered non-specific adaptation (adaptation), i.e. changes due to the main effect of "run". Third, a learning-specific change expressed on and only on the first day at run two was added to the adaptation effect (learning). Finally, a consolidation model (consolidation), in which learning-specific changes on the first day did not disappear but were consolidated – at half their level – by the second day was added to the adaptation effect. Practically, each of these four models

was specified with modulatory (bilinear) effects mediated by run-specific inputs that had different values between runs but were fixed over the duration of each run. These four profiles of coupling changes between runs were applied to different permutations of connections; namely, either forward or backward or both forward and backward between hippocampus and striatum. This produced ten unique models, because the three null models for different architectures were identical. This model comparison is quite subtle, in the sense that I tested for the presence or absence of changes in the context of full connectivity – not the presence or absence of connections per se.

Fulfilling my predictions of higher evidence for the learning or consolidation models, I found the highest log evidence for the consolidation model with a bidirectional change in connection strength between the hippocampus and striatum (see Figure 3-3a for an illustration). Remarkably, this was the winning model for 10 out of 11 participants (being the model with the second largest evidence for 1 participant; see Figure 3-3b).

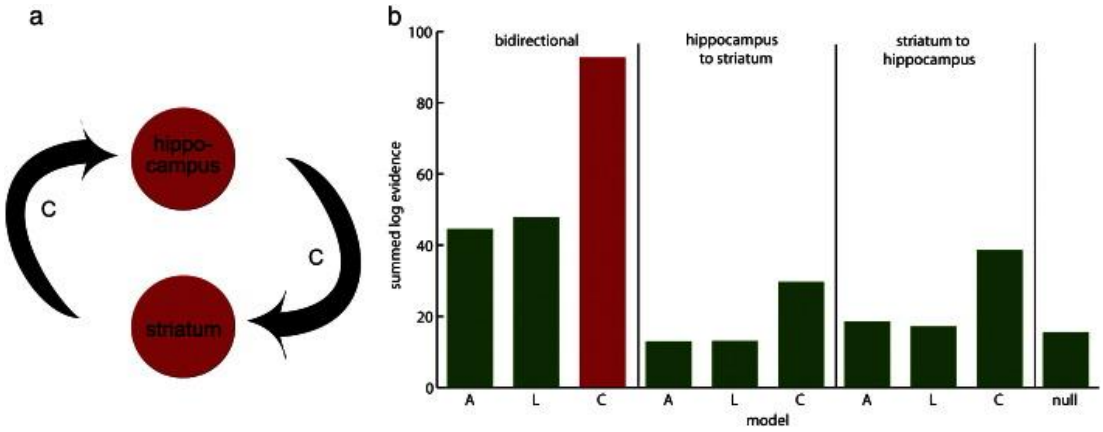

Figure 3-3 Winning model and summed (group) log evidence for all models. a) Schematic description of the model with a bidirectional connection between the hippocampus and the striatal region. The graphic shows which connections were modified by a consolidation pattern (see 3.3 for detailed explanation). b) The model plotted in a) showed the highest evidence (marked in red). Plotted is the summed log evidence per model relative to the model with the least evidence. The winning model was the same for almost all participants (10 out of 11). A = adaptation, L = learning, C = consolidation, null = no modification.

Having established the model with the highest evidence, quantitative changes in coupling were computed for each participant using a mixture of the run-specific changes as specified above (i.e. *adaptation* and *consolidation*) weighted by the appropriate run specific (bilinear) parameter estimates. These estimates (see Figure 3-4) provided a quantitative picture of the changes in coupling and its consistency over subjects. Reflecting the characteristics of the winning consolidation model effective connectivity changes were largest between the first and second rest run. They were smaller but consistent for the two rest runs on the second day of scanning, i.e. during rest runs three and four. The same pattern was observed for both directions, i.e. from hippocampus to striatum and vice versa.

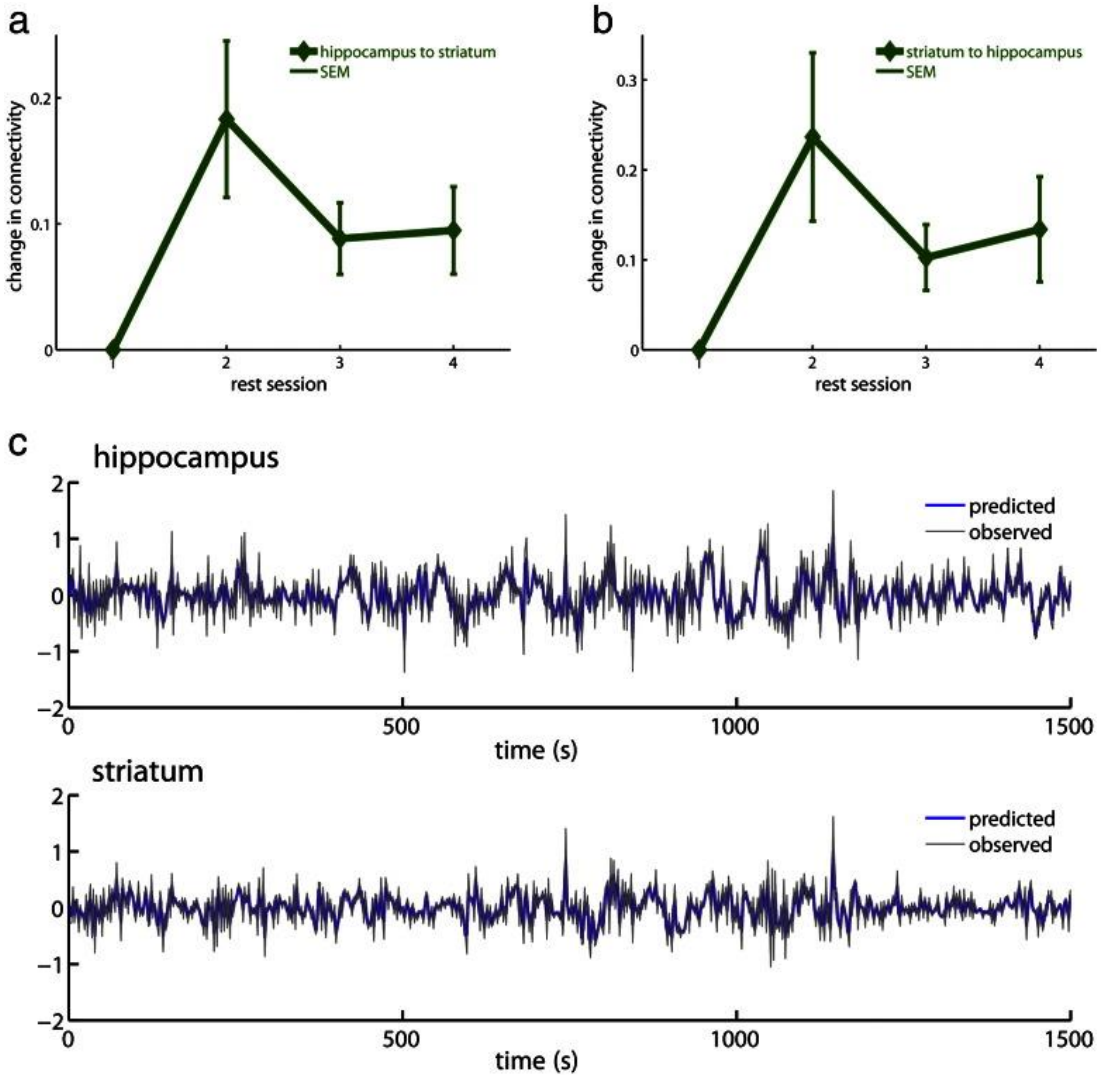

Figure 3-4 Parameter estimates and model fitting reflected consolidation. Parameter estimates for the modulation of the intrinsic connection from a) hippocampus to striatum and from b) striatum to hippocampus. After a big increase directly after the learning in rest run 2, the change in connectivity was preserved at a lower level on the second scanning day for both rest runs (i.e. rest runs 3 and 4). Plotted are the average values for all participants who learned the task (n = 11), error bars indicate the standard error of the mean (SEM). c) Overlay of observed (gray) BOLD time-series during rest with the time-series as predicted by DCM (blue). The two regions included in all tested models are shown for a representative participant.

3.4 Discussion

I investigated the neural correlates of the rapid perceptual learning phase in a standard visual paradigm and the relationship between learning related changes and

spontaneous fluctuations in resting state activity before and after that learning. I showed that a random dot coherence task can be learned by naïve participants within one training session. The task activated primarily visual and parietal brain areas. Significant learning related changes in neural responses were observed in the hippocampus. Furthermore, learning of the task had consequences for resting state connectivity: the hippocampal region changed its coupling with the striatum in a pattern that could be best explained in terms of consolidation. More precisely, a psychophysiological interaction analysis identified learning dependent changes in coupling with the hippocampus that were greater than equivalent changes on the second day without learning. Dynamic causal modelling of the directed interactions between the hippocampal and the striatal region showed that both forward and backward connections expressed learning dependent effects that persisted on the second day. This even allowed non-specific adaptation between paired runs on the two days of data acquisition.

While it is well known that performance on sensory tasks improves with practice, the time course of learning related changes is less established. Unlike mine, many studies do not investigate the early phase of learning, which is usually overlooked due to a familiarization period. This is particularly true for functional MRI studies. While some studies use difficult tasks with training over several days, weeks or even months (Blakemore and Campbell, 1969; Kahnt et al., 2011), rapid learning effects in a number of visual learning tasks have been reported after as few as 200 trials (Fiorentini and Berardi, 1981). Learning of a random dot coherence task, as used in this study, can occur after just 300 trials (Vaina et al., 1995). Using the same 2-alternative-forced-choice paradigm, participants improved their performance in a

single session from scoring close to chance to almost perfect. In a follow-up fMRI study Vaina et al. (1998) showed an increase in the activation in V5/MT and a decreased activation of the cerebellum, when comparing neuronal responses during the first task session with responses during the final session. However, the authors did not use any participant-specific performance measurement, whereas here I specifically identified participant-specific learning-related changes over time.

In line with several previous studies, my motion learning task activated visual areas including V5/MT (Newsome and Salzman, 1993; Rees, Friston, et al., 2000). The necessary role of the region for motion perception has been established in macaque monkeys and in human patients (Baker et al., 1991; Cowey and Marcar, 1992; Vaina et al., 2001), as well as in healthy humans using transcranial magnetic stimulation (TMS) (Walsh et al., 1999; Tadin et al., 2011). However, I did not find learning related changes (at the relatively conservative statistical threshold employed here) in any visual brain area. This might be due to the fact that my group of learners comprised only 11 participants. Thus, a potentially small effect in visual areas may not have been observed due to a lack of power. Importantly, my main interest here was not the specificity of the learning related effects, but the potential changes in connectivity during rest. For example, such changes are seen in a fronto-parietal network after participants learn a difficult shape identification task (Lewis et al., 2009).

My finding that early learning dependent effects were seen in the hippocampus supports the idea that sensory learning extends beyond a bottom-up process that is restricted to earlier sensory areas related to the representation of sensory stimuli.

Together with previous findings, my results suggest that different sensory learning tasks have different neural correlates in higher level brain areas. In line with several recent studies using electrophysiological and neuroimaging methods, my results are consistent with a role of non-sensory areas in visual decisions and learning (Law and Gold, 2009; Kahnt et al., 2011). Specifically, the role of the hippocampus as the classical area for explicit – or declarative – memory and spatial orientation has been challenged. For example, the MTL (including the hippocampus) is involved in tasks during which participants are not consciously aware of learned contingencies (Rose et al., 2011). Also, several hippocampal and parahippocampal regions including the subiculum change their activity in response to temporal regularities; demonstrating a role for human MTL in statistical learning and providing insight into the formation and evolution of memory representations (Schapiro et al., 2012).

The "classical" distinction between implicit and explicit learning is not straightforward for the motion tasks I used. Implicit learning refers to the incidental learning of complex information; i.e., without awareness of what has been learned (Sun et al., 2008). However, this definition is not uncontroversial (Frensch and Runger, 2003). Typically, three different stimuli structures are used to investigate implicit learning: patterns, sequences and functions (Forkstam and Petersson, 2005). In comparison, explicit learning has been characterised as a process similar to conscious problem-solving used for the control of task variables (Mathews et al., 1989), which gives rise to concrete and conscious knowledge about regularities in the environment (Reber, 1989). It is likely that the early learning phase of my task involved both types of learning. Indeed, the mechanism of any hippocampus-related learning processes does not appear to be sufficiently described by the established

dichotomy between explicit/implicit learning. On the one hand, hippocampal activity is associated with perceptual forms of associative learning (Fortin et al., 2002; Van Opstal et al., 2008); on the other, hippocampal involvement is seen for implicit higher-order sequence information (Lieberman et al., 2004), including visual sequence learning (Turk-Browne et al., 2010) and transitive inference tasks (Van Opstal et al., 2008). Furthermore, theoretical and empirical work has characterised the hippocampus as a fast learning system (Schendan et al., 2003; Colgin et al., 2008). I exposed my participants to only one learning session. The observed learning is thus classified as fast, compared to slow and usually small additional improvements over days, weeks or months.

While the traditional view of the role of the hippocampus has linked it to explicit/declarative learning (Penfield and Milner, 1958; Winocur, 1985; Neves et al., 2008) the striatum has been associated with implicit/non-declarative learning (Reiss et al., 2005; Wilkinson and Jahanshahi, 2007). However, as discussed, the classical dichotomy may no longer be tenable for the hippocampus, and may be obsolete for the striatum as well: first, my finding that the connectivity between the hippocampus and the striatum changes in response to learning during rest is in line with earlier findings suggesting that both the hippocampus and the striatum show a dynamic interaction during various types of learning (see Packard and Knowlton, 2002; Poldrack and Packard, 2003 for reviews). Moreover, several neuroimaging studies have examined the role of the hippocampus and the striatum during sequence learning using fMRI (Gheysen et al., 2011; Rose et al., 2011). These results highlight the importance of the MTL system and its connections with the striatum for perceptual learning, independent of its nature; i.e. implicit or explicit. My finding

that connectivity between the hippocampus and the striatum changes is particularly interesting with regard to their role in reinforcement learning. Reinforcement learning describes learning by trial and error to act in a way that maximises reward (Sutton and Barto, 1998). Previously, several studies have investigated the theoretical and empirical relation between perceptual learning and reinforcement signals (Seitz and Watanabe, 2005; Smith et al., 2009a). They showed that reinforcement learning can account for the learning during performance of a visual decision task (Law and Gold, 2009) driven by numerous cortical areas including the striatum (Schultz, 2007).

All these findings – including my own results – indicate that some learning related changes, and in particular early ones, involve non-sensory areas. These might involve an enhanced readout of sensory information as a result of behaviourally improved performance. In other words, fast learning may arise from changes in the interpretation of the respective sensory representation rather than changes in the sensory representation itself. More than that, the distinction between explicit and implicit learning systems seems to become more and more outdated (Rose et al., 2011).

From a methodological perspective, I present a practical example of the use of stochastic DCM for the analysis of fMRI resting state data. Li et al. (2011) established the validity of this method and its ability to model endogenous fluctuations in hidden neuronal states, thereby providing a new perspective on how regionally specific signals in fMRI are generated. Commonly used methods to investigate changes in connectivity are often based on correlations, thereby

addressing changes in so-called functional connectivity. However, functional connectivity does not support any conclusions about directionality, whereas DCM allows one to model (context dependent changes in) directed and possibly reciprocal connections between brain areas. In addition to deterministic, i.e. "classical" DCM, the newer stochastic DCM accommodates random fluctuations in hidden neuronal and physiological states. This approach may provide a more plausible perspective on how regionally specific signals in fMRI are generated.

3.5 Conclusion

In conclusion, I provide empirical evidence to show that the coupling of spontaneous fluctuations of a brain region engaged in early learning of a sensory task is changed during rest and that these changes persist for at least 24 h. Previously, it has been shown that task performance and/or learning leads to changes in the coupling between brain regions (Seitz et al., 2005; Stevens et al., 2010). Furthermore, performance in a novel perceptual task has been associated with the individual variability in functional connectivity during rest (Baldassarre et al., 2012). Here, I used recent advances in dynamic causal modelling to examine directed changes in brain connectivity in learning-related areas immediately and one day after learning. My key finding – that the coupling between a hippocampal and a striatal region are best explained by a consolidation model – provides further evidence for the idea that spontaneous fluctuations are continuously updated and modified by experience dependent plasticity. More generally, my findings support the view that the adult brain remains plastic throughout the life-span (May, 2011).

Chapter 4 The role of prestimulus activity in visual extinction

4.1 Introduction

4.1.1 The phenomenon of visual extinction

Visual extinction is commonly observed after right parietal damage. Patients with visual extinction perceive unilateral stimuli presented either in the left or the right visual field, but sometimes miss a stimulus in the left visual field during bilateral simultaneous presentation. Awareness of these left visual field stimuli is effectively "extinguished" by the stimulus in the right visual field. Visual extinction therefore offers a rare opportunity to study the neural correlates of perceptual awareness and unconscious processing.

4.1.2 How does visual extinction relate to spatial neglect?

The nosology of visual extinction is not clear. It could either represent a component, or mild form, of the classical visuospatial neglect syndrome (Vallar, 1993; Heilman et al., 1994; Rafal, 1994) or a completely different type of visuospatial attention deficit (Umarova et al., 2011). Some data suggest a dissociation between the two syndromes (Vallar et al., 1994; Hillis et al., 2006; Vossel et al., 2011), whereas others emphasise the similarity, especially when the lesions are clustered in the inferior parietal lobule (Posner et al., 1984; Vallar et al., 1994; Rees, Wojciulik, et al., 2000; Vuilleumier et al., 2010). Umarova et al. (2011) compared the activation patterns of acute stroke patients with neglect and visual extinction during visuospatial processing and found an increased activation in the left prefrontal cortex only for patients with extinction. These results suggest that visual extinction and

neglect are separate syndromes. However, this study used only unilateral stimuli and did not identify the areas involved in the extinction of the left stimulus during bilateral stimulation. Interestingly, the right inferior parietal cortex has been implicated in the simultaneous processing of bilateral targets (in animal studies (Lynch and Mclaren, 1989) and healthy participants (Ciçek et al., 2007).

4.1.3 Mechanisms of visual extinction

Several previous studies have investigated the neural mechanisms of visual extinction, using bilateral and unilateral stimuli. Essentially, two different approaches have been employed. The first approach investigates residual cortical processing of the extinguished stimulus by comparing responses in bilateral extinguished trials with responses in unilateral right trials; i.e., trials with different physical properties that lead to the same behavioural response. Contrasting these experimental conditions using fMRI shows that the extinguished stimulus in the left visual field activates early visual cortex, as well as the extrastriate visual cortex in the damaged right hemisphere, e.g. (Vuilleumier and Rafal, 2000; Driver et al., 2001; Rees, Kreiman, et al., 2002; Rees, Wojciulik, et al., 2002; Vuilleumier et al., 2002). A cross modal study using the same paradigm with tactile information reported activation of primary sensory cortex (S1) in response to extinguished stimuli (Sarri et al., 2006). These results provide a potential explanation for the unconscious processing assessed using indirect measures such as priming, e.g. (Baylis et al., 1984; Berti et al., 1987; Ladavas et al., 1993; Vuilleumier and Rafal, 2000; Driver et al., 2001; Vuilleumier et al., 2002).

The second approach examines the neural correlates of awareness by comparing seen and unseen stimuli during bilateral presentation; i.e., trials with the same physical properties leading to different behavioural responses. Converging evidence from several studies supports the idea that the interplay between posterior visual areas and fronto-parietal circuits is crucial for a visual stimulus to reach awareness, e.g. (Driver et al., 2001; Rees, Kreiman, et al., 2002; Rees, Wojciulik, et al., 2002). Thus, it has been suggested that a pathological bias in attention towards the ipsilesional visual field leads to the "extinction" of the contralesional stimulus from awareness during bilateral stimulation. This is in line with the observation that the colour and form of the extinguished stimulus can still be processed to a certain extent. In short, the parietal damage might compromise spatial awareness and responding, rather than disrupting early visual processing.

4.1.4 Prestimulus activity affects perception

As outlined in Chapter 1, it is well known that ongoing or intrinsic neuronal activity influences subsequent evoked responses. Furthermore, prestimulus activity has been related to systematic variations in behaviour and thus is functionally significant. For example, Fox et al. (2007) found that 74% of spontaneous trial-to-trial variability in button press force can be accounted for by ongoing fluctuations in the intrinsic activity in somatosensory cortex. Similarly, correlations between ongoing fluctuations of brain activity and perception are observed across different paradigms and different species (Ress et al., 2000; Giesbrecht et al., 2006; Hesselmann, Kell, and Kleinschmidt, 2008; Hesselmann, Kell, Eger, et al., 2008). Fluctuations in prestimulus activity in visual areas measured with EEG and MEG influences the

detection of upcoming stimuli (Mathewson et al., 2009; Wyart and Tallon-Baudry, 2009). Specifically, alpha activity in somatosensory areas might play a crucial role in optimising neuronal processing, thereby influencing behaviour (Haegens et al., 2011). In addition, fMRI results suggest that the BOLD signal in a cortical area preferentially responding to faces is higher preceding experimental trials that are perceived as faces compared to vases using an ambiguous figure (Hesselmann, Kell, Eger, et al., 2008). In motion coherence tasks, BOLD signals in motion-responsive brain areas are higher before trials that are perceived as showing coherent compared to random motion (Hesselmann, Kell, and Kleinschmidt, 2008). Finally, a recent fMRI study extended the investigation of fluctuations in ongoing brain activity to the domain of cognitive control: prestimulus activity in several task relevant regions – including higher cognitive areas – scales with the size of the Stroop effect (Coste et al., 2011). In sum, there is strong evidence that endogenous variations in prestimulus neuronal activity bias subsequent perceptual decisions.

4.1.5 Can I analyse visual extinction using prestimulus activity?

Here, I set out to answer the question of how it is possible that patients with visual extinction sometimes see and sometimes miss the left stimulus during bilateral stimulation. My strategy was to compare prestimulus BOLD signals before bilateral visual stimulus presentation depending on whether the trial was subsequently categorised as a "bilateral seen" or as a "bilateral unseen" trial; in other words, whether the patient failed to detect the stimulus in the left visual field. I focused on visually response areas and used a simple detection paradigm with bilateral and unilateral face stimuli. First, I identified visually responsive areas in a patient

showing visual extinction. Second, I compared prestimulus activity in these regions during bilateral stimulation with and without extinction. Finally, I used DCM (Friston et al., 2003) to examine whether changes in the coupling or excitability of these regions could explain both prestimulus activity and subsequent differences in stimulus bound responses. Specifically, I investigated whether extinction might be mediated by a difference in intrinsic (within area), or extrinsic (between areas), effective connectivity, i.e. the causal influences that neural units exert over one another (Friston, 1994), or sensitivity to neuronal afferents. DCM is the method of choice for my question because it tests hypotheses or models that are cast in terms of directed connections among neuronal populations. As outlined in 2.3.6, this contrasts with less informed approaches – such as functional connectivity – that simply measure (undirected) correlations between haemodynamic responses at different points in the brain.

4.2 Materials and methods

4.2.1 Participant

One male patient (IPJ) aged 66 with left visual extinction (following a right parietal stroke, see Figure 4-2) gave informed consent to participate in the study. The participant showed left visuospatial neglect on four standard clinical measures – see 4.2.2.1. Functional imaging was conducted 3 years and 4 months post-stroke. IPJ was suited for in-depth study as he had a structurally intact visual cortex in the right hemisphere, despite suffering from enduring visual extinction on clinical confrontation and formal computerised testing. However, he showed lower left

quadrant visual field impairment. Therefore, all experimental stimuli were presented in the upper visual quadrants.

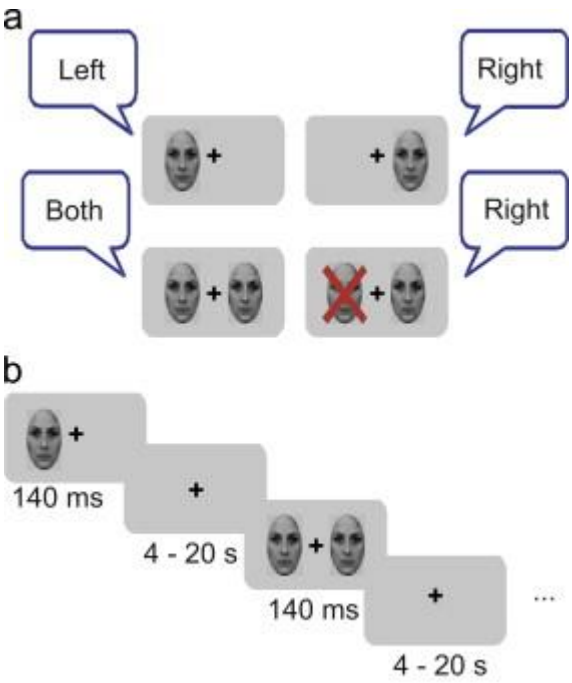

Figure 4-1 The extinction paradigm. (a) Facial stimuli were presented unilaterally in either the left or right visual fields (upper row) or bilaterally (lower row). Depending on the response of the patient, trials were categorised after scanning as bilateral seen (BS) (lower row, left) and bilateral unseen (BU) (lower row, right) trials. (b) Stimuli were presented for 140 ms (or 120 ms respectively during the later blocks due to learning effects of the patient) and were segregated by an intertrial interval ranging between 4 and 20 s.

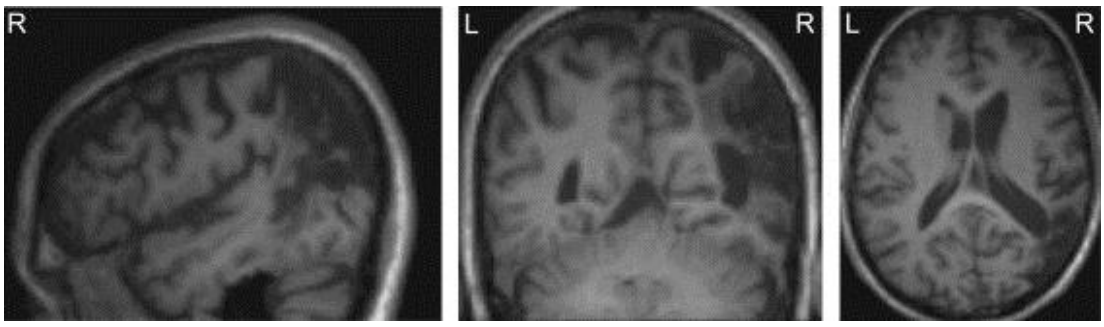

Figure 4-2 Right parietal lesion. T1-weighted structural MRI scan acquired during the first of two scanning sessions where the pre-existing right parietal lesion is clearly apparent.

4.2.2 Design and procedure

The experiment was approved by the local ethics committee.

4.2.2.1 Neuropsychological testing

Prior to functional imaging, IPJ was tested for clinical signs of visual extinction by confrontation. In addition, he was presented with bilateral, unilateral left and unilateral right visual stimuli outside of the scanner to titrate the different parameters for the scanning sessions. The criteria defined by Vallar et al. (1994) (i.e. >30% misses of left events during bilateral stimulation, but <20% misses of single left events during unilateral stimulation) were used. In addition, IPJ performed three standard clinical neuropsychological measures to test for signs of visuospatial neglect: in the cancellation task, IPJ was presented with an A4 sheet of paper containing circles and crosses. Half of these contained a small gap, which had to be crossed out (15 on each side, i.e. 30 in total). Typically, patients with neglect fail to cancel targets located on the left side of the page.

During the line bisection task, e.g. Wilson et al. (1987), IPJ was presented with three 18 cm lines printed in the middle of A4 sheets of paper and was asked to put a mark where he thought the middle of each line was. Neglect patients tend to underestimate the leftmost side of the line, thus making errors by deviating rightwards from the true midpoint. In the lateral preference task, which measures spontaneous lateral attentional biases, the patient was shown 10 pairs of virtually arranged, identical, left-right mirror-reversed chimeric face stimuli – joining together left and right halves of the same face posing different neutral or happy expressions. The patient was asked to judge whether the upper or bottom face looked happier. Right

hemisphere damaged patients with left neglect typically select the face that is smiling on the right side of the display, e.g. Sarri et al. (2010), which is opposite of that for healthy participants, e.g. Mattingley et al. (1993).

4.2.2.2 fMRI paradigms

After the behavioural data had been analysed, IPJ was tested during two scanning paradigms using fMRI (on separate days), which I refer to as the "extinction paradigm" and the "stimulus localiser". During both paradigms he was asked to fixate centrally and to respond with the right hand on a keypad.

4.2.2.2.1 Extinction paradigm (event related design)

Each trial of the extinction paradigm comprised the presentation of faces on the left, the right or both sides. Stimuli were presented for 140 ms (run 1–6) or 120 ms (run 7–9). The duration was shortened during the last three runs to ensure an equal number of bilateral seen (BS) and bilateral unseen trials (BU), as IPJ improved in terms of visual detection. His task was to indicate where he saw a stimulus or stimuli respectively. The conditions were presented in random order and the inter-trial interval was randomised to minimise anticipation, ranging between 4 and 20 s. See Figure 4-1 for a visual description of the extinction paradigm. Each run comprised 35 trials, with 23 bilateral stimulus presentations and 12 unilateral presentations (i.e. six, for each side). IPJ completed nine runs divided over two scanning sessions (five in the first session), resulting in 207 bilateral and 54 trials for each side respectively.

4.2.2.2.2 Stimulus localiser (block design)

Each trial of the stimulus localiser entailed the presentation of faces, objects or scrambled images on the left or right side. To elicit detectable responses in visual areas, stimuli were presented for 250 ms with an inter-stimulus interval of 500 ms. Based on previous experiments on visual extinction, I used longer stimulus presentation times (compared to the extinction paradigm) thereby increasing the efficiency or sensitivity of detecting visually responsive areas in the lesioned brain. IPJ completed two runs each consisting of 12 blocks during which each stimulus (faces left or right, objects left or right, scrambled images left or right) appeared twice (i.e. 12 trials per block). Each block was followed by a 6 s break (i.e. blank screen). To ensure fixation throughout, the task was to press a button whenever the fixation cross turned red. Note that, unlike the event related extinction paradigm, this paradigm was a more efficient and longer block design that only presented unilateral visual stimuli. This enabled us to identify visually responsive areas for subsequent analysis in an efficient way. Furthermore, because I was particularly interested in the mediation of extinction in early visual cortex (providing ascending sensory information to higher category selective regions), I averaged over all stimuli types in the localiser to define functionally preserved visual responses at lower levels in the visual hierarchy.

4.2.2.2.3 Stimuli

All stimuli were presented at the same location in the upper quadrants of the visual field, subtending 4.91 × 6.70 of visual angle. Face stimuli were taken from a face database provided by the Karolinska Institute, Stockholm, Sweden (Oosterhof and Todorov, 2008), and were cropped, resized, and presented in greyscale. Scrambled

images were derived from the object and face images via a random exchange of picture elements organised in a 20×20 matrix.

4.2.3 fMRI data acquisition

A 3T Trio MRI scanner (Siemens Medical Systems, Erlangen, Germany) with a standard head coil was used to acquire functional data with a standard echo planar imaging (EPI) sequence (matrix size 64×64; field of view 192×192 mm; in plane resolution 3×3 mm; 32 slices in descending acquisition order; slice thickness 3 mm; echo time 30 ms; TR 2 s). IPJ attended two scanning sessions separated by 1 week. During both sessions, fieldmaps were acquired to correct for geometric distortions in the EPI images due to inhomogeneities of the magnetic field. Finally, a structural T1-weighted scan was acquired during each session (field of view 256×240 mm; in-plane resolution 1×1 mm; 176 sagittal slices of thickness 1 mm; echo time 2.98 ms). Each run included five dummy volumes that were discarded during the data analysis to allow for T1 equilibration.

4.2.4 Data analysis

4.2.4.1 Behavioural data

Data from the extinction paradigm were analysed with regard to correct trials and reaction times. These were compared among different conditions using repeated measures ANOVAs.

4.2.4.2 fMRI data

Functional data were analysed using SMP12 (http://www.fil.ion.ucl.ac.uk/spm/software/spm12/). Pre-processing of the data

involved realignment of each scan to the first scan of each run, coregistration of the functional data to the structural data of each day and, finally, coregistration of the structural scan of the second day, to co-register all the functional images. The functional data were smoothed with an 8 mm Gaussian kernel after spatial normalisation to the MNI template brain. The data were filtered with a standard 128-s cut-off, high-pass filter to remove low-frequency drifts (including differences between runs), while preserving as much variance due to spontaneous fMRI fluctuations as possible (Cordes et al., 2001). Statistical tests were family wise error rate corrected (FWE) for multiple comparisons at $p < 0.05$ or uncorrected at $p < 0.001$ across the entire brain.

4.2.4.2.1 Extinction paradigm

The time-series of each functional run were analysed using a standard general linear model (GLM) including eight regressors for the four conditions or trial types of interest: right and left unilateral trials and bilateral trials on which the stimulus was seen (BS) or unseen (BU): each condition had two regressors, one for the prestimulus baseline and one for the stimulus evoked responses. The prestimulus baseline was modelled as a 6 s long period starting 7 s before stimulus onset (allowing a 1 s gap between baseline and stimulus presentation). The choice of 7 s was based upon informal model comparisons, using models of sustained prestimulus activity starting 3 s and 5 s before stimulus onset (not reported) and heuristics based upon the timescale of fluctuations in resting state fMRI studies. These fluctuations have a characteristic length of about 10 s, which places an upper bound on the duration of sustained endogenous activity.

The evoked responses were modelled as standard event-related stick functions. Note that the prestimulus baseline and event related response regressors for each trial type were necessaryly correlated, because one precedes the other. However, because the haemodynamic response function peaks between 4 and 6 s, the activity modelled by the two regressors could be estimated with reasonable efficiency. I did not orthogonalise these regressors, which means that any significant prestimulus baseline effects discovered cannot be explained by event related differences.

Stimulus functions were convolved with a canonical haemodynamic response function to provide regressors for a standard general linear model (GLM). Movement parameters in the three directions of motion and three degrees of rotation were included as confounding regressors of no interest. Contrasts of parameters of the effect of interests were estimated over all nine task runs. The associated statistical parametric maps (SPMs) were used to test for differences in the neural activity during the prestimulus period of BS and BU trials.

4.2.4.2.2 Stimulus localiser

The time-series of both functional runs were analysed with a standard GLM comprising six regressors modelling the effects of faces, objects and scrambled images for left and right side, using even-related regressors. Again, movement parameters were included as confounding regressors of no interest. Contrasts of parameters were estimated over both task runs. The resulting SPMs were used to test for differences in the neural responses between right and left visual field stimulation to identify regions showing visual responses to lateralised stimuli.

4.2.4.2.3 Peristimulus time histograms (PSTH)

To quantify the time course of the BOLD activity in the regions of interest (ROI) showing differences between seen and unseen trials (i.e. BS vs. BU), I estimated event related responses in these ROIs using a finite impulse response (FIR) convolution model. The parameters of the corresponding GLM report BOLD activity in successive time bins of 2 s of peristimulus time (in my case). I evaluated event related responses over all nine runs from 7 s before to 9 s after stimulus presentation.

4.2.4.2.4 Dynamic causal modelling (DCM)

The standard SPM analyses described above localised (visually responsive) brain areas that showed higher activity before BS compared to BU trials. My hypothesis was that perception depends upon prestimulus baseline activity and that this activity depends upon fluctuations in extrinsic or intrinsic connectivity. In the final analyses, I used DCM to determine whether differences in connectivity between seen and unseen trials were intrinsic to the visual regions showing prestimulus baseline effects and/or in the extrinsic connections between these regions.

My comparisons of effective connectivity were based on Bayesian model comparison using deterministic DCM (Penny et al., 2004). To test for differences in effective connectivity I concatenated the data of all nine runs and used three regressors: one for the prestimulus baseline of all bilateral trials (using 7 s boxcar functions: the duration of the prestimulus period was extended to 7 s, to ensure that prestimulus conditions were maintained until the stimulus arrived), one for the stimulus onset of all bilateral trials (using a standard event related stimulus function)

and one to indicate whether the stimulus was seen or not (i.e. using the same boxcar function as for the first regressor but only for BS trials).

I created 16 models corresponding to a 4 × 4 factorial design with two factors. All models included reciprocal extrinsic connections between the two visual areas of interest (the areas are referred to as "right" and "left" subsequently), which were driven by the prestimulus and stimulus onset effects described in Sections 4.2.2.2.1 and 4.2.2.2.2. The first factor was extinction-dependent differences in intrinsic connections of the two regions (with the four levels: both, left, right, or neither), while the second factor was differences in extrinsic connections between those two regions (with the four levels: both, left-to-right, right-to-left, or neither). Crucially, both the prestimulus and stimulus related driving effects were identical for seen and unseen trials. The only difference between seen and unseen trials was mediated by a prestimulus effect that modulated connections within (intrinsic) or between (extrinsic) the two regions. In other words, the extinction of the left stimulus could only be explained by a difference in (intrinsic or extrinsic) connectivity or sensitivity to presynaptic inputs that was established before the arrival of the stimulus.

All 16 models were fitted to the concatenated time series of the extinction runs using standard variational Bayesian model inversion. The relative evidence for each model was approximated with variational free energy to provide the posterior probability of each model (assuming uniform priors over subsets of families of models that were compared) (Friston et al., 2003; Stephan et al., 2009). I used a two-step heuristic search for the best model: First, I assessed the contribution of changes in intrinsic connectivity by assessing the posterior probability for the four different families of

intrinsic connection strength changes (effectively averaging over my uncertainty about putative changes in extrinsic connections). I then compared the four different extrinsic models within the winning intrinsic family.

Finally, I examined the modulation of connections, i.e. changes in connection strength, using the parameter estimates for the intrinsic and extrinsic connections of the winning model. Note that in this DCM, the modulatory or bilinear effects are modelled by additive changes to the connection strengths. This means that the modulatory values alongside the connections in Figure 4-7 should be added to the coupling parameters associated with each connection. The ensuing modulation of connections by a prestimulus effect presupposes an endogenous fluctuation in the local synaptic processes that determine effective connectivity. In other words, the prestimulus effect is an effect on coupling strength (quantified by DCM) that causes changes in neuronal activity (quantified by SPM).

4.3 Results

4.3.1 Patient showed signs of visual extinction

Four typical clinical neuropsychological measures of neglect were used to test for signs of visual extinction. In the cancellation task IPJ missed three targets on the left side and none on the right side. In the line bisection task, IPJ's mean deviation error toward the right when indicating the middle of the line was 3.3 cm. In the lateral preference task the patient chose faces with the smile on the right side in nine out of 10 cases. During confrontation IPJ missed the stimulus presented in his left visual

field in nine out of 10 bilateral trials. He did not miss any of the unilateral left trials. Thus, he fulfilled the criteria defined by Vallar et al. (1994).

4.3.2 Stimulus localiser activated visual areas

Comparing BOLD signals for stimulus presentation in the left visual field (i.e. independent of stimulus type) to those for presentation in the right visual field, I found activations in three regions in the right hemisphere (see Table 4-1), including primary visual areas and precuneus. The opposite contrast, testing for regions that were more active during presentation of a stimulus in the right visual field, revealed activation of left primary visual cortex. However, the activation was much more confined. See Figure 4-3 and Table 4-1 for detailed results.

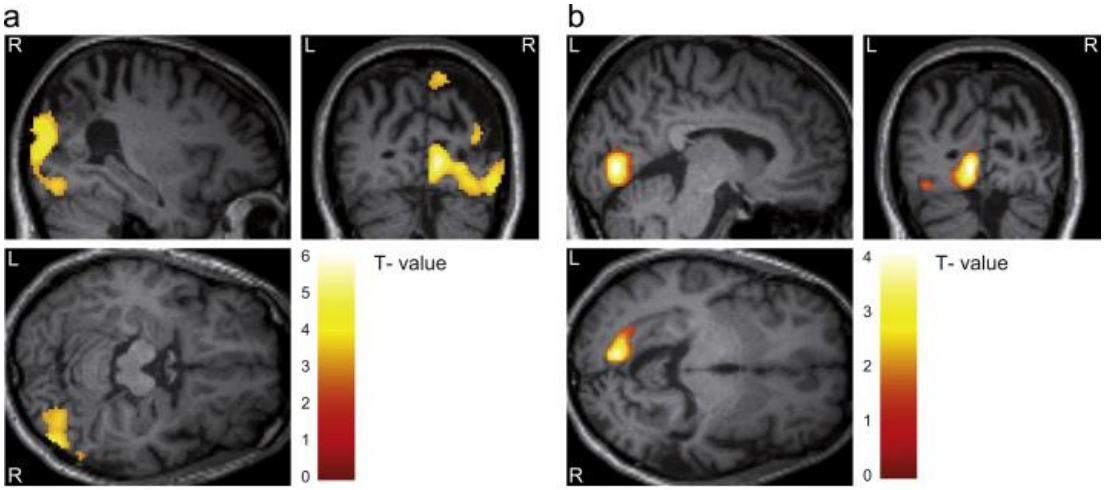

Figure 4-3 Stimulus localiser activated visual areas in both hemispheres. (a) Activations due to stimulation of the left visual field were confined to the right hemisphere. Images are displayed at $p < 0.001$, uncorrected for illustration purpose. (b) Activations due to stimulation of the right visual field were confined to the left hemisphere and showed much less distributed pattern. Images are displayed at $p < 0.01$, uncorrected for illustration purpose.

Table 4-1 Stimulus localiser activated visual areas. Activations during the stimulus localiser left > right are restricted to the right hemisphere, and vice versa. Directions refer to visual fields. [a]Voxel-level statistics, $p < 0.05$, FWE. [b]Voxel-level statistics, $p < 0.001$, uncorrected and a clustersize of at least 10 voxels. L = left hemisphere, R = right hemisphere, left = left visual field, right = right visual field.

	MNI coordinates				
Left > right	x	y	y	t-value	P-value
R BA 17 (including calcarine sulcus)	12	-82	0	6.17	< 0.0001[a]
R BA 19 / occipital medial	34	-84	14	5.03	= 0.011[a]
R Precuneus	10	-74	60	3.94	< 0.0001[b]
R Inferior orbital frontal	54	44	-12	3.29	= 0.001[b]
R Superior occipital	28	-82	46	3.27	= 0.001[b]
Right > left					
L BA 17 (including calcarine sulcus)	-8	-80	-2	4.16	< 0.0001[b]

4.3.3 Extinction paradigm produced unseen trials

Averaged over the nine runs of the extinction paradigm, IPJ missed 45% of bilateral trials (corresponding to 94 out of 207 trials) – these are the BU trials. There was no significant difference between BS and BU trials over the nine runs ($F_{1,8} = .97$, $p = .35$). He reported seeing 94 % (50 out of 54 trials) of unilateral left trials, and 98 % (53 out of 54 trials) of unilateral right trials. The difference in seen unilateral trials was not significant over the nine runs ($F_{1,8} = 4.00$, $p = 0.08$). Average response times for the BS trials were longest, with unilateral trials being faster than bilateral trials; however, reaction times did not differ significantly between the different trials ($F_{3,24} = 1.70$, $p = .20$). See Figure 4-4 for the details of the responses and reaction times.

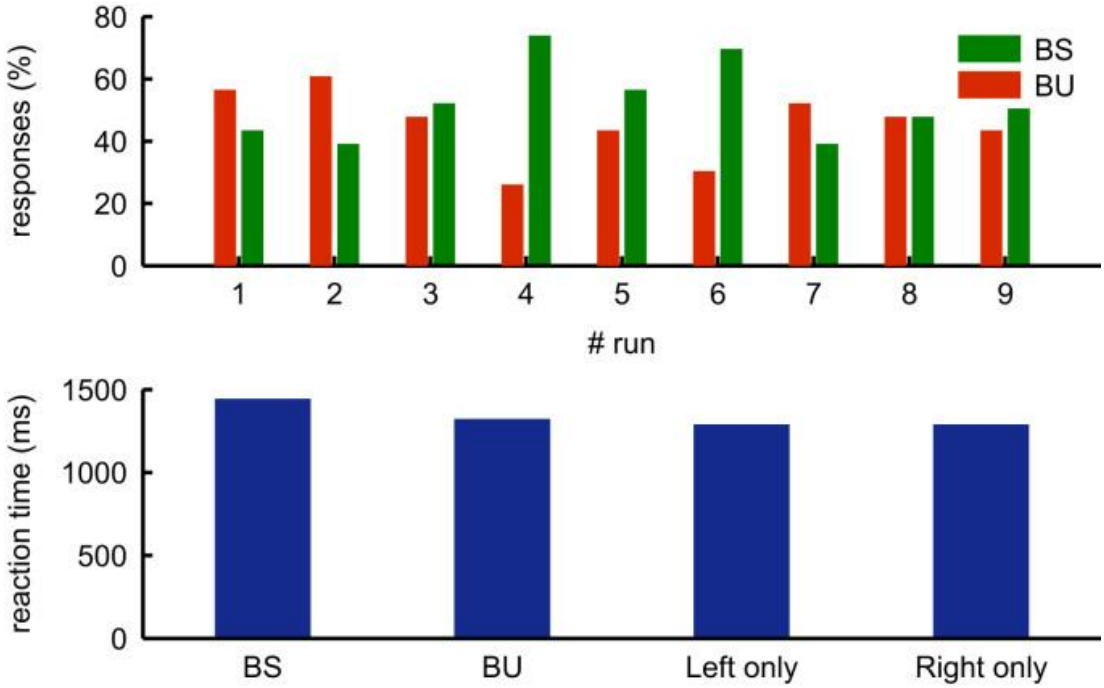

Figure 4-4 Behavioural results of the extinction paradigm. Top: Percentages of BS and BU trials for all nine runs. Bottom: Reaction times for the four trial types (see text) averaged over all nine runs.

4.3.4 Prestimulus activity in visually responsive areas affects perception

I identified regions showing higher activity before BS compared to BU trials by comparing the BOLD signal between these two conditions in a 6 s prestimulus baseline window starting 7 s before stimulus presentation. Crucially, I found an overlap with visual areas that were activated by the stimulus localiser in both hemispheres: BA 19/ occipital inferior right cortex (MNI x = 36, y = −78, z = −16, t = 3.32, $p < 0.001$ uncorrected) and BA 17/ calcarine sulcus left (MNI x = −4, y = −86, z = −8, t = 3.28, $p < 0.001$ uncorrected). The overlap between the visual responses to bilateral stimuli and the localiser stimuli was substantial: 82 % (65 out of 79 voxels) of the BS–BU activation in the right hemisphere overlapped with the activation due to the stimulus localiser in the right hemisphere, 47 % (69 out of 147

voxels) of the BS–BU activation in the left hemisphere overlapped with the activation due to the stimulus localiser in the left hemisphere. See Figure 4-5.

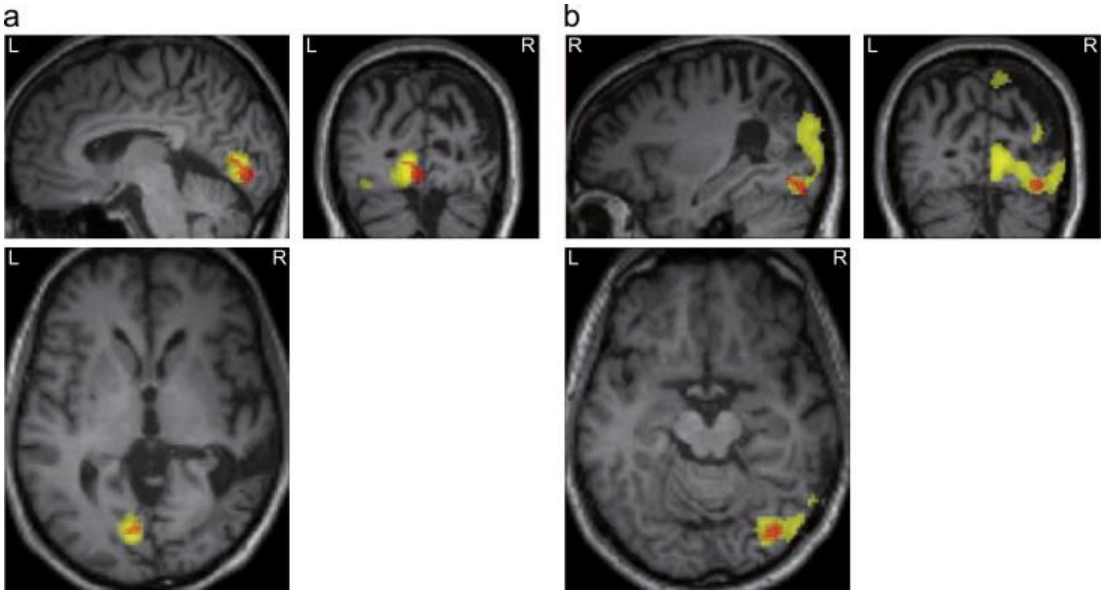

Figure 4-5 Visually responsive areas are more active before bilateral seen trials. Overlay between lateralised visually evoked responses during the stimulus localiser (yellow) and the prestimulus baseline effects revealed by the extinction paradigm (red), showing the areas that are more active before BS compared to BU trials. (a) Left hemisphere. (b) Right hemisphere. All overlays are displayed at $p < 0.005$, uncorrected for illustration purpose.

In addition, an exploratory analysis (using an uncorrected threshold of $p < 0.001$) revealed several regions showing an effect in the same direction, including activity differences in the brain stem and parietal cortex. See Table 4-2 for an overview. The opposite contrast, i.e. higher activity before BU vs. BS trials, revealed no region that would survive FWE correction. The closest was a right inferior frontal area (MNI $x = 54$, $y = 16$, $z = 4$, $t = 3.43$, $p < 0.001$, uncorrected).

Table 4-2 Activity differences during the extinction paradigm for the baseline period testing for areas that show higher activity before BS compared to BU trials. Voxel-level statistics at $p < 0.001$ uncorrected.

	MNI coordinates				
	x	y	y	t-value	P-value
Brainstem right	16	-30	-20	3.54	< 0.0001
SMA right	4	-22	58	3.39	< 0.0001
Paracentral lobule right	12	-40	52	3.33	< 0.0001
Parahippocampal region right	14	-6	-22	3.21	= 0.001
SMA left	-4	8	54	3.18	= 0.001
Brainstem left	-10	-28	-22	3.16	= 0.001
Rectus left	-10	46	-20	3.11	= 0.001
Frontal medial left	-44	56	18	3.11	= 0.001

4.3.5 Time-course of responses to seen and unseen trials

To quantify the prestimulus fluctuations in BOLD responses, I used an FIR model for responses in the two visual areas that showed increased activity before BS compared to BU trials. Both ROIs show a distinct increase in their haemodynamic response before stimulus presentation, which starts to diverge between seen and unseen trials as early as 5 s before stimulus onset. See Figure 4-6 for details.

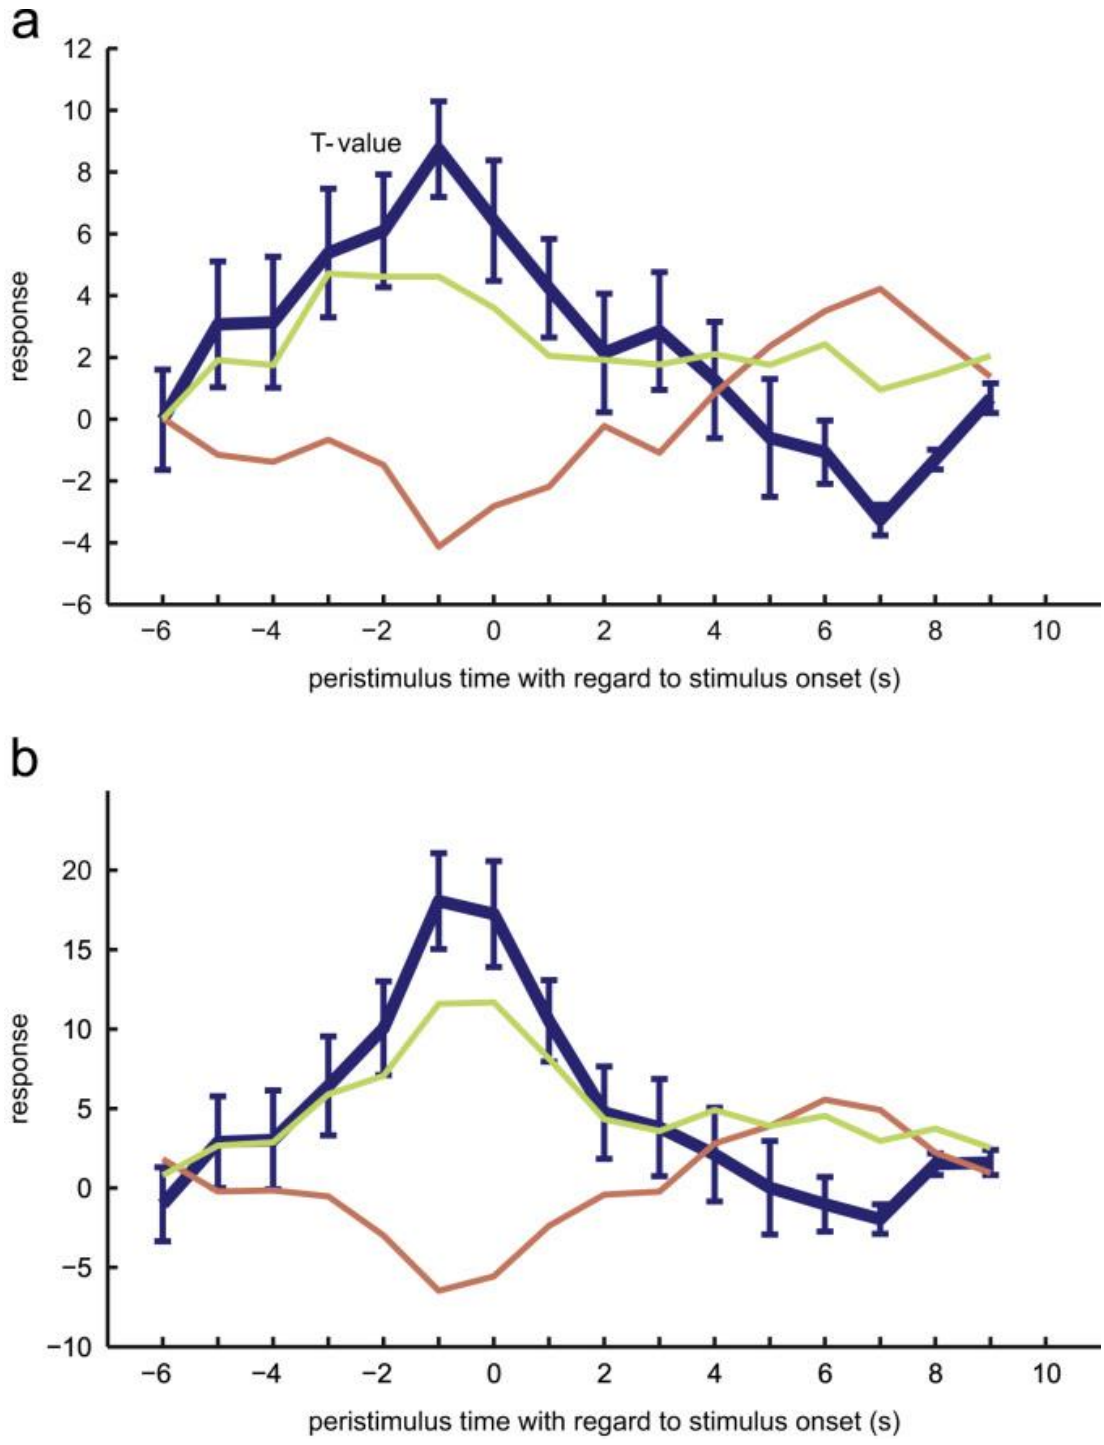

Figure 4-6 Peristimulus time courses show difference before stimulus onset. The time courses of the two visually responsive areas showed a differential activity during baseline before stimulus presentation depending on whether the subsequent bilateral stimulus is seen. Plotted are the difference with SD (blue) and the time course for BS (green) and BU (red) individually. (a) ROI in right occipital inferior cortex. (b) ROI in left calcarine sulcus.

4.3.6 Perception depends on the coupling between visual areas

Having identified two visually responsive areas that showed increased activity preceding BS trials, compared to BU trials, I next asked whether the connectivity within and between those two regions differed before stimulus exposure. The models tested differed in terms of where differences in connectivity were expressed depending on whether a bilateral trial was seen or not. Sixteen models as described in 4.2.4.2.4 were fitted and compared in terms of the posterior probabilities. The first comparison between intrinsic families showed that I could be 99 % confident that there was an effect on intrinsic connections and 73 % confident that both visual areas were involved (although there was a 26 % probability that only the left area was affected). Following this comparison, I compared the four models within the winning intrinsic family (were both intrinsic connections changed). This comparison showed that I could be 99 % sure that there was a change in extrinsic connections and 68 % confident that both efferent and afferent connections to the lesioned hemisphere were involved (although there was a 30 % chance that just the right to left extrinsic projection changed).

Having selected the most plausible model, I looked at the differences in effective connection strength between seen and unseen trials. For BS trials, effective connectivity within and between the two areas increased. See Figure 4-7 for details. Crucially, all intrinsic and extrinsic effective connection strengths were elevated prior to seen trials. For the intrinsic connections, this entailed a decrease in self-inhibition, between 60 % (on the left) and 20 % (on the right). The remarkable thing about the changes in extrinsic connectivity is that they (both) change from being

mildly inhibitory to being excitatory. Quantitatively, these changes were more marked in the right-to-left extrinsic connection. In short, changes in both intrinsic (decreased self-inhibition) and extrinsic (from mildly inhibitory to excitatory) appear to precede stimuli that are subsequently seen.

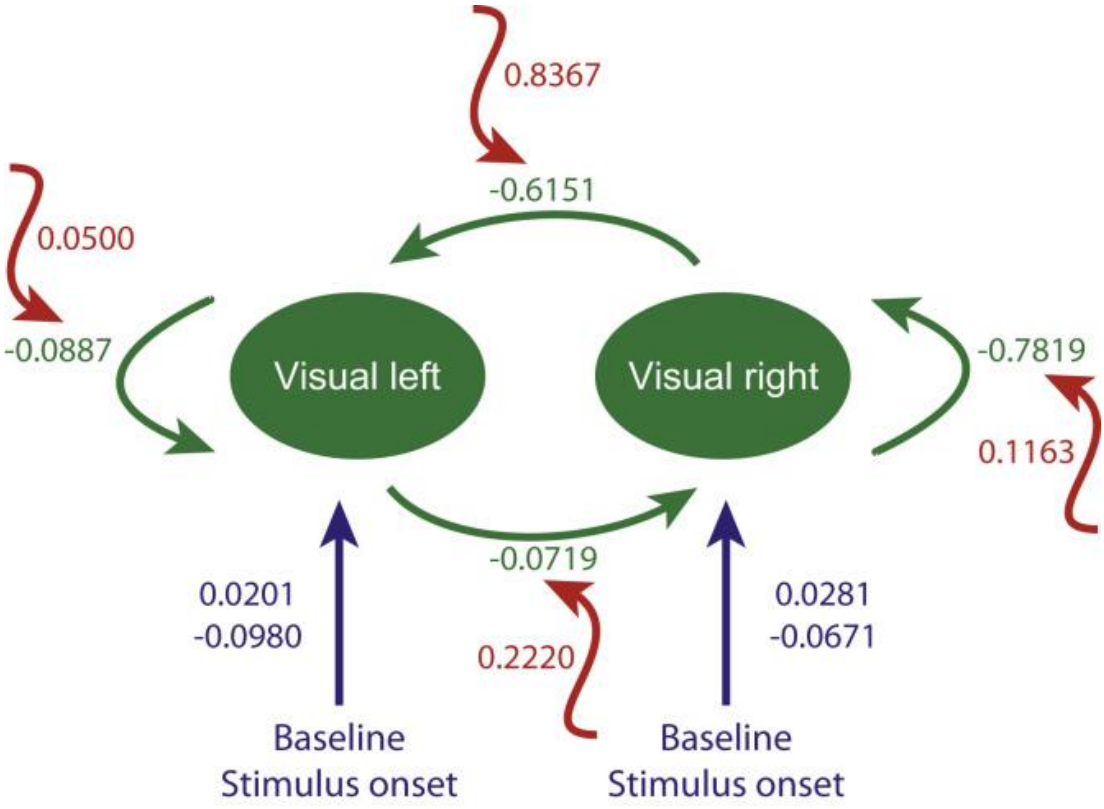

Figure 4-7 Differences in effective connectivity before bilateral seen trials. The winning model including its inputs (blue), intrinsic and extrinsic connectivity (green) and modulations (red) is shown. The numbers describe the parameter weights. Effective connectivity was increased for all four intrinsic and extrinsic connections prior to bilateral seen trials. Visual left = calcarine sulcus, visual right = inferior occipital cortex.

4.4 Discussion

The aim of this case study was to address two questions: Does prestimulus activity in visually responsive areas in a patient with visual extinction predict subsequent perception (as seen in healthy subjects in other tasks), and do fluctuations in

connectivity between these regions determine neuronal and perceptual responses? I used a simple detection paradigm with unilateral and bilateral phase presentation. I concluded that fluctuations in connectivity between regions that exhibited higher activity prior to bilateral seen compared to bilateral unseen trials provide a sufficient account of both baseline fluctuations and perceptual reports. This finding is consistent with studies of normal subjects. However, care should be taken in generalising this conclusion to the normal brain. This reflects the Catch-22 associated with lesion-deficit studies: I can only study the correlates of extinction in the lesioned brain, which means that I cannot exclude the possibility that the physiological (fluctuating connectivity) basis of neuronal and perceptual responses is itself pathological. Having said this, one could argue that the consilience between my results and studies of baseline fluctuations in normal subjects (Fox and Raichle, 2007; Hesselmann, Kell, and Kleinschmidt, 2008; Hesselmann, Kell, Eger, et al., 2008) suggests one might find the same changes in connectivity, were it possible to study perceptual extinction in the healthy brain.

4.4.1 Prestimulus activity in visual areas affects stimulus perception

Our results are in line with previous work on visual extinction and the visual areas identified by these. In fact, the two areas that show a higher prestimulus activity prior to bilateral seen trials are very close to the visual areas reported by Rees et al. (2002), when investigating the unconscious residual cortical processing of the extinguished stimulus in the contralesional visual field. I extend the results of previous studies showing that visual areas can be activated without leading to awareness, e.g. Sarri et

al. (2010) by providing evidence for the idea that activity prior to stimulus presentation is indicative for subsequent perception.

Furthermore, the activations in response to unilateral trials in the present study were very similar to the regions reported by Rees et al. (2002) for the same contrast: in both cases, responses in the lesioned right hemisphere were greater and more widespread than in the left (see Figure 4-3).

4.4.2 Prestimulus activity in other brain areas might play a role

In addition to the two visual areas, I identified several brain regions that showed signalling differences that were associated with subsequent conscious perception during bilateral stimulation; however, these failed to survive correction for a whole brain search, i.e. they did not survive FWE correction (possibly reflecting the relatively low efficiency of my single case study). Among these areas are frontal and parietal regions, which have been identified in previous studies of visual extinction (see below). In fact, during stimulus processing the interplay between posterior visual areas such as the ones found here and a fronto-parietal network seems to be crucial for perceptual awareness, e.g. Vuilleumier and Rafal (2000); Driver and Vuilleumier (2001); Rees, Kreiman, et al. (2002); Rees, Wojciulik, et al. (2002). In addition, I detected higher prestimulus activity prior to seen bilateral trials bilaterally in the brainstem. This evolutionary old part of the brain is known to control autonomic functions of the peripheral nervous system and modulate arousal and alertness, two criteria that may be important in determining awareness. Indeed, alertness levels are known to modulate the severity of spatial neglect (Robertson et al., 1998; Malhotra et al., 2006; George et al., 2008) and low alertness has even been

linked with neglect-like rightward biases in healthy participants (Manly et al., 2005) including in extinction tasks (Matthias et al., 2009).

4.4.3 Mechanisms behind visual extinction

I used Bayesian model comparison to investigate potential changes in the coupling between the two visually responsive areas identified prior to the stimulus. I found the highest probability for models that allowed an increase in both intrinsic and both extrinsic connectivity for sensitivity preceding bilateral stimuli that are subsequently seen. In case of the intrinsic connections these changes represented a decrease in self-inhibition. Remarkably, the extrinsic connections changed from being mildly inhibitory to being excitatory. It should be noted, that real extrinsic connections between the two areas are both excitatory (using the neurotransmitter glutamate). However, in DCM, effective connections can be polysynaptic and an extrinsic connection can be effectively inhibitory (presumably by targeting inhibitory interneurons). Quantitatively, the changes in effective connectivity were more marked in the right-to-left extrinsic connection, i.e. from the lesioned to the healthy hemisphere. Crucially, these changes in connectivity for sensitivity were sufficient to explain both the differences in baseline activity prior to stimulus onset and the perception dependent differences in stimulus bound responses.

These results suggest that fluctuations in cortical gain or excitability (both to intrinsic and extrinsic presynaptic inputs) may underlie the decreased neuronal response and a failure to perceive stimuli that are subject to extinction. This is interesting in that exactly the same mechanisms – at the synaptic level – are thought to underlie attentional modulation, which may be dysfunctional in extinction. Furthermore, they

speak to the precision-dependent explanation for detecting signals based upon predictive coding; in the sense that precision is thought to be encoded by postsynaptic gain (Feldman and Friston, 2010) and that optimising postsynaptic gain corresponds to attention. This is important because the many mechanisms modulating postsynaptic gain include the classical modulatory neurotransmitter systems, originating in the brainstem (see above). A heuristic (and overly simplistic) explanation for these results could be as follows: spontaneous fluctuations in ascending aminergic and cholinergic neurotransmitter systems result in spontaneous fluctuations in the effective gain of neuronal populations in visual cortex, both to intrinsic and extrinsic afferents. If the resulting increases precede a stimulus, then the neuronal responses evoked by stimulus are amplified and gain access to higher hierarchical levels, enabling deeper processing and perceptual inference – and subsequent perception.

4.4.4 Limitations of the study

In this work, my primary focus was on early visual mechanisms that might underlie fluctuations in the perception of stimuli. From this perspective, the current case study represents a lesion-deficit model that enables the comparison of seen and unseen stimuli and their physiological correlates. Generalising my conclusions – about the underlying role of intrinsic and extrinsic connectivity – to the normal brain clearly rests on the assumption that both the perceptual and physiological processing of seen and unseen stimuli are quantitatively the same in my patient and the normal population.

One might also argue that my findings would be more plausibly generalised if I had been able to reproduce the results using further patients with extinction. This is certainly the case and extinction has a reasonably high prevalence. However, despite testing several patients with extinction, only the patient reported here was considered suitable for fMRI. Although this is a single case study, one can be reassured by the fact that fMRI produces an enormous amount of data and the degrees of freedom I have used for my analyses were much greater than any conventional group study. Having said this, this case study should probably be regarded as proof of principle, until reproduced in other people.

4.4.5 Methodological aspects

From a methodological perspective, I present a practical example of the use of DCM in a patient with a parietal lesion. Frequently used methods to investigate changes in connectivity are often based on correlations and address changes in so-called functional connectivity, which describes statistical dependencies between spatially segregated neuronal events. However, this approach does not support any conclusions about directionality or the distinction between intrinsic and extrinsic influences. In contrast, effective connectivity is based on a mechanistic model of how the observed data were caused and allows the modelling of directed and reciprocal connections within and between brain areas.

Although DCM is an established procedure; for those people less familiar with the analysis of fMRI time-series, DCM can be contrasted with alternative procedures: in general terms, distributed interactions, as measured by fMRI, can be characterised in terms of either functional or effective connectivity (see 2.3.5 and 2.3.6). Functional

connectivity refers to the statistical dependence or correlations between observed responses (Biswal et al., 1995; Cordes et al., 2001), while effective connectivity refers to the underlying and directed connections strengths that cause correlations (Friston, 1994). Analyses of effective connectivity generally use DCM, although other techniques have been tried (such as structural equation modelling (SEM), multivariate/vector autoregressive models (MAR/VAR) and Granger causality, see 2.3.6.1). DCM is unique in that it incorporates an explicit model of neuronal interactions and allows for region specific neurovascular coupling. If these regional differences are ignored, they can lead to false inferences about effective connectivity (David et al., 2008). DCM is therefore the only approach that allows one to test hypotheses about connectivity at the neuronal level. More precisely, it uses a neurobiologically plausible model of neural population dynamics and a biophysically plausible forward model which describes the transformation from neural activity to the measured hemodynamic signal (Goebel et al., 2003; Stephan and Friston, 2011). Consequently, it is possible to fit the parameters of the neural and the forward model in a way that predicted time series are optimally similar to the observed ones.

4.5 Conclusion

In conclusion, I studied a patient with visual extinction after a right parietal lesion that spared visual cortex. I was able to extend previous work showing that activations in visual, parietal and frontal areas can be observed without awareness, e.g. Sarri et al. (2010). In doing so, I have tried to infer the mechanisms that determine whether extinction will occur during bilateral stimulation. I found that the prestimulus activity in two visual areas in both hemispheres showed increased activity prior to bilateral

seen stimuli compared to those that were unseen. In addition, I used DCM to examine directed changes in coupling within and between these two areas and found that all four intrinsic and extrinsic connections were increased for several seconds prior to stimulus onset. In line with previous studies of prestimulus activity and its role in perception, my results support the idea that prestimulus activity in distinct brain areas is an important determinant of subsequent perception and behaviour.

Chapter 5 Effects of ongoing cortical state on ambiguous perception

5.1 Introduction

The perception of visual images is not a one-to-one mapping of visual inputs to particular perceptual representations. Even with exactly the same visual input, one can perceive different structures at different times. For instance, the well-known Rubin face/vase stimulus demonstrates how perceived shape and meaning can vary across time without a change in the incoming sensory information (Rubin, 1915). Other well-known examples of this perceptual mutli-stability are the Necker cube (Necker, 1832) and the rabbit-duck illusion (McManus et al., 2010) These phenomena suggest that the visual system does not operate purely reflexively. Instead, whether I perceive a particular interpretation of a stimulus partly depends on the current state of the brain at the time sensory information is received. Spontaneous fluctuations in brain activity and their functional role have been investigated using a wide range of neurophysiological methods (Arieli et al., 1996; Kenet et al., 2003; Holcman and Tsodyks, 2006; Fox and Raichle, 2007, see 1.2). For instance, work using electrophysiological recordings and neuroimaging has suggested that the often unexplained variance in spontaneous activity contributes significantly to evoked responses elicited by external sensory stimuli (e. g. Fox et al., 2007; Deco and Romo, 2008; Sadaghiani et al., 2009; Mennes et al., 2010).

One paradigm that has been used by visual neuroscientists to understand the role of ongoing neuronal fluctuations is the so-called random dot motion coherence task.

The stimuli in this task typically involve a set of dots moving within an aperture, see Figure 5-1 and Chapter 3.

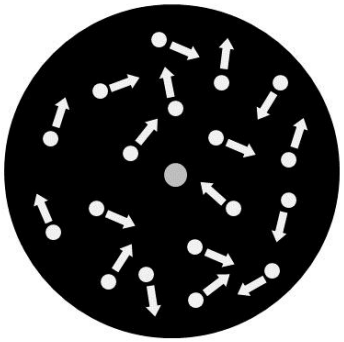

Figure 5-1 Random dot motion stimulus. Dots are moving in a circular aperture, indicated by the arrows. Most of the dots are moving in a random direction, while a certain low percentage is moving in the same – coherent – direction.

Some proportion of the dots move coherently in the same direction whereas others move randomly in other directions. When a relatively large percentage (in this study over 30 %) of the dots moves coherently in one direction, participants invariably see the stimulus as having a coherent motion in a particular direction. With an intermediate level of dot motion coherence, however, participants can be equally likely to see either coherent motion or non-coherent motion even with exactly the same coherence level on different trials. Hesselmann et al. (2008) used fMRI to show that prestimulus baseline activity in the motion-sensitive occipito-temporal cortex (hMT+) predicts whether participants perceive a random dot stimulus presented at perceptual threshold as coherent or random. Higher prestimulus baseline activity in hMT+ just before the stimulus biased participants to see an ambiguous stimulus as containing coherent motion. However, fMRI has a low temporal resolution and thus does not reflect higher frequency fluctuations in activity that may be related to this effect. In comparison, EEG and MEG have a high temporal resolution and have been

used in similar paradigms to investigate the role of prestimulus activity of different frequency bands on perceptual outcomes (e.g. Hanslmayr et al., 2007; Hipp et al., 2011). In particular, Donner et al. (2007, 2009) found that motor-responses selective MEG activity in the beta and gamma frequency ranges predicted participants' reported perceptual outcome in in a comparable design several seconds before the actual behavioural choice was made by participants. More precisely, they observed a gradual built up of this choice predictive activity during stimulus viewing. However, in contrast to other studies using similar perceptual paradigms (e.g. Linkenkaer-Hansen et al., 2004), they did not find any performance-predictive activity in the prestimulus interval comparing errors with correct behavioural choices (Donner et al., 2007). This might be due to the fact that their experimental setup was designed such that the participants' baseline state was maximally controlled, thereby minimising the contribution of ongoing fluctuations in brain activity and processing capacity. First, trials were presented rapidly after each other. Second, stimuli were presented for a long duration (2 to 3 s), also reducing the impact of occasional differences in ongoing activity before stimulus onset on perceptual judgments. Third, participants were not instructed to attend to a specific location prior to stimulus onset inducing preparatory activity in the dorsal pathway (Corbetta et al., 2002; Sapir et al., 2005). In contrast, I was specifically interested in the possible influence of slow baseline fluctuations. Therefore, I chose different settings, using a long interstimulus interval, a very short presentation time and instructed participants to fixate centrally throughout the experiment.

Crucially, the observed effects mentioned above using functional MRI, EEG and MEG are purely correlational. To test for causality, one could use transcranial

magnetic stimulation (TMS) that uses specific frequencies to stimulate particular brain regions. Prior to such an experiment, it needs to be determined which frequencies are important for the perceptual outcome.

Here, I set out to test whether specific prestimulus signatures of intrinsic oscillatory activity predicted perceptual outcome in a random dot motion (RDM) coherence task. Following Hesselmann et al. (2008) I used a temporally sparse paradigm in which I presented RDM stimuli every 5 to 7 seconds while recording scalp EEG. By sparsely arranging stimuli with a random gap between them, I attempted to decouple each stimulus from any systematic effects of the immediately preceding stimulus and any systematic temporal expectations. Furthermore, by allowing long periods between stimuli this allowed for non-stimulus, spontaneous fluctuations in activity to occur (and be measurable).

To capture changes in spontaneous, ongoing electrical brain activity, I subjected the prestimulus EEG data on each trial to spectral decomposition to estimate the power in various frequency bands. By averaging prestimulus power (in each frequency) separately for trials subsequently perceived as coherent and those perceived as containing random motion and then comparing these, I could determine whether the oscillatory activity in different frequency bands was different just before the two different perceptual outcomes. With the high temporal resolution of EEG, I was able to estimate the precise timing of the relevant fluctuations as well as characterise the type of activity (e.g. frequency band, phase lag) which is not possible with functional MRI.

I used post-stimulus physiological activity to validate participants' subjective reports of stimulus coherence. Participants had to respond to three types of stimuli with different coherence levels, namely with high coherence that was easily detectable, with no coherence, and with threshold coherence based on the individually determined coherence level that elicited a coherent percept in 50 % of the cases. I first compared the event-related potentials (ERPs) elicited by physically different stimuli types, i.e. trials with a high coherence versus trials with no coherence, trying to replicate the findings of Niedeggen and Wist (1999). I predicted that the same differences should be observed comparing the ERPs elicited by physically identical trials with a coherence level at individual threshold that elicit distinctive responses, i.e. comparing random versus coherent periliminal trials.

In short, my post-stimulus measures corroborated my participants' subjective reports of how they perceived the ambiguous stimulus on each trial. My analysis of pre-stimulus oscillations on each trial will assess whether spontaneous fluctuations in ongoing brain activity are correlated with the reported perceptual outcome of ambiguous stimulus processing. I will assess this across multiple frequency bands of the EEG up to a second preceding the stimulus onset.

5.2 Materials and methods

5.2.1 Participants and apparatus

Fifteen healthy right-handed participants (six female, average age of 27.8 years, range 21 – 36 years) took part in the EEG experiment. They gave written informed consent prior to the experiment and received a remuneration of £ 10 per hour. The

experiment was approved by the UCL Research Ethics Committee (Project ID: 1161/001). Participants were seated in a dark, shielded room with their heads on a chinrest. Stimuli were presented on a 1280 x 1024 CRT monitor at a viewing distance of 65 cm.

5.2.1.1 Stimuli

Stimuli were dynamic dot displays of 500 white squares (size 0.2°) randomly distributed on a dark grey circle (23°). Participants were instructed to fixate on a central white square (size 1°) throughout the experiment. For 355 ms, stimuli moved up- or downward at 14° / s, and with different coherence levels. Noise dots moved in a "random walk". Signal dots had a limited lifetime of six frames. Participants were asked to report as quickly and accurately as possible after stimulus presentation. Via a button press they indicated whether they had perceived coherent or random motion. Responses were given with the left and right hand for either percept; the assignment was counter-balanced across participants.

5.2.1.2 Training and thresholding

During a piloting study I observed significant inter-individual differences in initial task performance. Therefore, I developed a training practice tailored to initial performance of every participant which was performed by each participant before the EEG data was collected. Every participant completed short blocks of 40 trials (ISI 2 s), half of which were coherent. The coherence level of coherent trials started at 40 % for the first block(s) until the participant achieved at least 80 % correct. Subsequently, the coherence level was reduced by 5 % or 10 %. This change was determined by the experimenter's assessment of their performance with better

performance leading to the 10 % decrease. As with the first block, this one was run until 80 % performance was reached. This procedure was continued until the participant reached 80 % correct for a block with a coherence level of 20 %. After each trial, written feedback was given on the screen. The number of necessary training blocks varied between three and ten blocks between participants.

After the training, I used the method of constant stimuli (30 trials of 7 different coherence levels i.e. 2, 6, 10, 14, 18, 22, and 26 %, presentation order randomised, no feedback, ISI 2 s) to determine individual motion coherence thresholds for each participant individually (50 % level of a cumulative normal distribution fit; average motion coherence threshold across participants 13.7 %, range 6 – 19 %).

5.2.1.3 Behavioural task during EEG

During the EEG recording, three motion coherence levels were used, subliminal (0 % coherence, 20 trials), periliminal (individually estimated threshold, average coherence of 13.7 %, 60 trials), and supraliminal (30 % coherence, 20 trials). Between stimuli, the fixation square remained on the screen. The ISI was 5 – 7 s given by a uniform distribution, to avoid expectancy effects and allow for the signal to stabilise after the motor response and possible eye blinks. All but one participant performed four blocks of the task (one did three only due to time constraints). Trials were presented in a pseudo-random order such that each coherence condition preceded equally often each of the three coherence levels, ruling out unbalanced carry-over effects from previous stimuli within each participant's trial order (Brooks, 2012). This constraint resulted in blocks of 98 or 99 trials; however, the exact

amount of analysed trials inevitably varied between participants due to differences in response patterns and the rejection of artefacts during EEG data pre-processing.

5.2.1.4 EEG data acquisition

We used a BioSemi ActiveTwo system with 64 active electrodes cap (BioSemi, Amsterdam, The Netherlands) in conjunction with LabView BioSemi software to acquire the data at a sampling rate of 1024 Hz. There were no filters applied at recording other than the digitisation rate. The Common Mode Sense active electrode and the Driven Right Leg passive electrode formed the ground during recording and these were positioned to the left and right of POz, respectively. In addition, I recorded from the nose tip, mastoids, as well as the horizontal and vertical eye electrodes.

5.2.1.5 fMRI data acquisition and analysis

A structural MRI scan and a functional localiser for hMT were acquired from all participants using a 3T Siemens Trio MRI scanner (Siemens Medical Systems, Erlangen, Germany) with a standard 32 channel head coil. For the localiser a standard echoplanar imaging (EPI) sequence (matrix size 64 x 64; FOW 192 x 192 mm; in plane resolution 3 x 3 mm; 32 slices in ascending acquisition order; echo time 30ms; acquisition timer per slice 69 ms; TR 2.176 s) was used. A structural T1-weighted scan was acquired on the same day (matrix size 256 × 240; field of view 256 × 240 mm; in-plane resolution 1 mm × 1 mm; 176 sagittal slices of thickness 1 mm; echo time 2.48 ms; acquisition time per slice 7.92 ms).

During the functional localiser, participants looked at both static and moving dots (towards or away from the central fixation point) at a low contrast. Individual hMT

were identified for each participant using a contrast comparing moving dots with static dots. Average coordinates, in MNI space, for let hMT were (-42, -75, 5) and for right hMT were (46, -68, 2). See Table 5-1 for individual coordinates. All functional images were normalised to MNI space for further processing.

Table 5-1 Individual MT coordinates. Average and individual MNI coordinates of left and right MT for each participant. S9 did not obtain an MRI session and average coordinates were used.

Left MT	X	Y	Z	Right MT	X	Y	Z
Average	-42	-75	5	Average	46	-68	2
S1	-48	-72	-2	S50	42	-68	-4
S2	-36	-70	4	S52	45	-64	7
S3	-42	-68	16	S53	50	-66	8
S4	-44	-80	8	S54	52	-66	2
S5	-42	-73	7	S55	51	-67	4
S6	-42	-74	8	S56	50	-64	4
S7	-40	-70	12	S57	46	-66	4
S8	-42	-76	4	S58	52	-72	0
S10	-38	-82	-4	S60	40	-74	0
S11	-38	-82	4	S61	46	-74	-2
S12	-48	-66	6	S62	44	-64	-4
S13	-44	-74	0	S63	46	-68	4
S14	-36	-82	-2	S64	38	-70	-2

5.2.2 EEG data analysis

Data were pre-processed in SPM8 (http://www.fil.ion.ucl.ac.uk/spm/software/spm8/, Wellcome Trust Centre for Neuroimaging, London, UK) and statistical analysis was conducted in Fieldtrip (Oostenveld et al., 2011).

5.2.2.1 Pre-processing

EEG sensor locations were determined using either the default SPM values (seven participants) or using individual electrode and fiducial locations obtained in Brainsight Neuronavigation (Rogue Research, Montréal) when available (eight participants) and then co-registered with the anatomical MRI scan. Data were down-sampled to 512 Hz and re-referenced to the average of all electrodes. "Bad" channels were defined as those with persistent artefacts across approximately more than 20 % of trials as identified manually by the authors. To avoid any "bad" channels contaminating the average reference, bad channels were identified before re-referencing. This resulted in excluding eleven channels across four subjects. Eye-blinks were identified by SPM's automated blink detection algorithm based on activity in the VEOG electrode (placed below right eye). Epochs were created around these blinks (-200 ms to 200 ms) and these were passed to a singular value decomposition (SVD) to identify principal components associated with blink artefacts. Based on the spatial topography of the component, the first one or two components were associated with blink artefacts. Sensor data was corrected for the detected artefacts by removing the blink-associated component(s) according to the Berg method (Berg and Scherg, 1994). The corrected time series was then epoched from -1500 ms to 500 ms around stimulus onset and baseline corrected. Removal of any remaining artefacts (e.g. muscle, etc.) was performed manually using the FieldTrip Visual Artefact Rejection tool by removing channels and trials with variance appeared to be outliers.

5.2.2.2 ERP analysis

Event-related potentials (ERPs) in the post-stimulus data were calculated for two different comparisons by standard averaging within each of the four conditions for each subject: (1) supraliminal coherent versus subliminal random and (2) periliminal coherent versus periliminal random trials. After averaging, each ERP was low-pass filtered at 30 Hz and baseline corrected to the interval from -200 ms to 0 ms, i.e. stimulus onset. Because of a priori observations from previous studies about motion-associated ERP components (Niedeggen and Wist, 1999; Kuba et al., 2007), comparisons between conditions were restricted to an a priori spatio-temporal region-of-interest (ROI) comprising posterior electrodes ('P1', 'P2', 'P3', 'P4', 'P5', 'P6', 'P7', 'P8', 'P9', 'Pz', 'PO3', 'PO4', 'PO7', 'PO8', 'POz', 'O1', 'O2', 'Oz') and post-stimulus times 100 ms to 500 ms. Inferential statistics for each comparison were conducted across the entire ROI using a dependent-samples cluster-based t-test procedure (Maris and Oostenveld, 2007). The cluster-based tests used a cluster-forming threshold of $p = 0.05$; 20,000 Monte Carlo permutations; and a weighted cluster mean statistic (weight = 1, emphasising peak intensity) (Hayasaka and Nichols, 2004). This procedure provides weak control of family-wise error rate (FWER) across the ROI with no assumptions about the auto-correlations in the data (Groppe et al., 2011).

5.2.2.3 Prestimulus analysis

Prestimulus spectral power was computed, for each trial, in two broad bands with parameters specific to each band. Low frequency (2 – 30 Hz) power was estimated from -1.0 s to -0.1 s using a 500 ms Hanning window in 12 ms steps and a frequency resolution of 2 Hz. High frequencies (25 – 100 Hz) were analysed using the SPM

multitaper method with a 5 Hz frequency resolution, a window length of 400 ms, time steps of 12 ms and a time bandwidth of 7. Raw resulting power values were used. The resulting trial-by-trial time-frequency maps were averaged across trials within each condition and participant.

5.3 Results

5.3.1 Behavioural results

On average, participants perceived periliminal trials as coherent in 51 % of trials (range 35 – 70 %, STD 12.33). Subliminal trials were correctly perceived as random in 91 % of trials (range 69 – 100 %, STD 8.98), resulting in an average false alarm rate of 9 %. Supraliminal trials were correctly perceived as coherent in 88 % of trials (range 62 – 100 %, STD 10.63). The differences were assessed in a one-way ANOVA ($F_{2,39} = 60.04$, $p < .001$). Subsequent post-hoc testing (using Tukey's test) revealed that periliminal trials were significantly different from subliminal ($p < 0.001$) and supraliminal trials ($p < 0.001$); subliminal and supraliminal trials did not differ significantly ($p = 0.76$). The response pattern is shown in Figure 5-2 and indicates that participants perceived values at their individual threshold as coherent approximately half of the time.

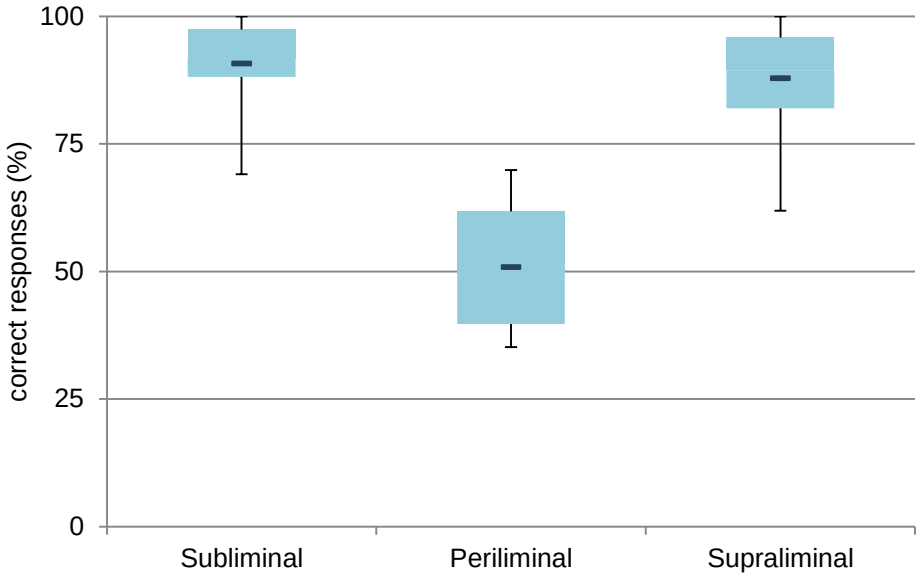

Figure 5-2 Response pattern across participants. Percentages of correct responses are given for subliminal, periliminal and supraliminal trails indicated by the bold vertical bar (darker blue). Whiskers indicate the first and third quartiles of the distribution. Although the response bias in some participants led to variability in the range of periliminal trials perceived as coherent or random, all participants showed a high sensitivity in their detection.

Using supraliminal and subliminal trials, *D'* was calculated for each participant and ranged between 1.88 and 4.07 (average 2.78), reflecting good performance for trials were the signal was clearly present (supraliminal) or absent (subliminal) (MacMillan and Creelman, 2005). The response bias *beta* ranged from .23 to 6.8 (average 2.2), reflecting the variability in response patterns between participants; *beta* is a measure of the willingness of responding 'coherent' for a certain trial and therefore indicates how much evidence is needed before participants report 'coherent' for any given trial. For all participants, arbitrary responses or guessing could be excluded based on adequate performance in both subliminal and supraliminal conditions.

5.3.1.1 Performance and response patterns

Performance over the time course of the experiment was constant. For each of the blocks, the performance over time did not change ($F_{11,3}$ = .242, p = .867) and there

was no interaction between condition and block number ($F_{11,6}$ = .579, p = .747). The response pattern over blocks is shown in Figure 5-3. A decrease in performance due to fatigue or an increase due to learning of the task can thus be ruled out.

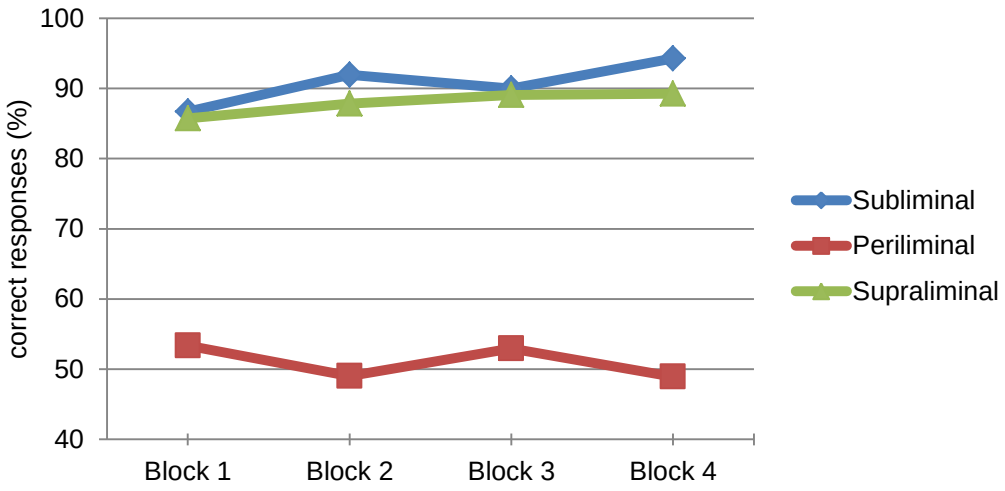

Figure 5-3 Response pattern over blocks. Response pattern are averaged over participants. The percentage of correct responses remained the same over the course of the experiment for all trial types, indicating that no effects of fatigue or learning were observed.

To test for possible carry-over effects of previous percepts across all trials and participants, I looked at the pattern of response repetitions (i.e. the probability of consecutive "coherent" or "random" responses). My results are in line with those reported by Hesselmann et al. (2008): the repetition pattern approximated a geometric distribution (goodness-of-fit $R2$ = 0.97 for coherent percepts, $R2$ = 0.98 for random percepts, see Figure 5-4). This indicates that behavioural reports were stochastic in nature and that no carry-over effect from previous trials was observed.

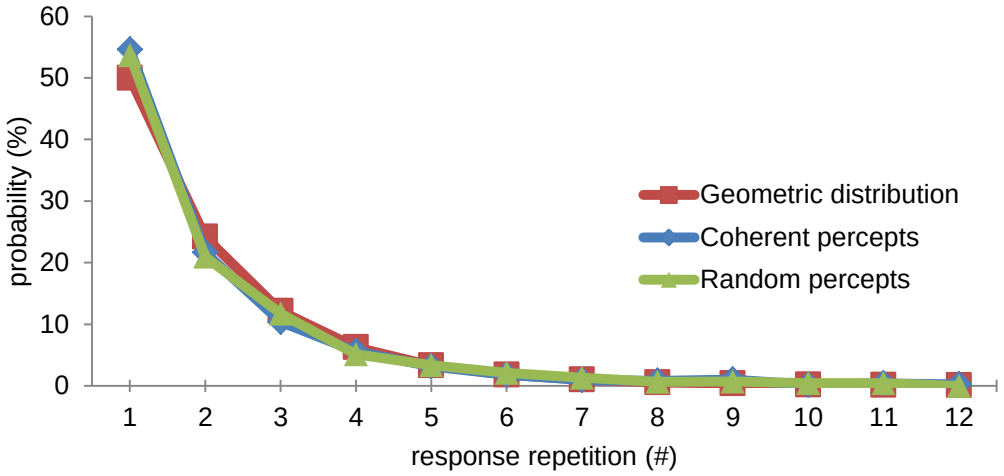

Figure 5-4 Response repetitions for random and coherent percepts. Response repetitions approximated a geometric distribution. "Coherent" and "random" responses showed a high goodness-of-fit, indicating no carry-over effect of percepts from previous trials.

5.3.1.2 Reaction times

Median response was 651 ms for supraliminal trials, 655 ms for subliminal trials and 664 ms for periliminal trials ($F_{2,39} = 0.74$, $p = .48$), indicating that participants did not need longer to respond to a particular condition. Next, I compared correct and erroneous responses separately using a 2-way ANOVA. Neither stimulus type ($F_{2,74} = 2.63$, $p = 0.17$), response outcome ($F_{2,74} = 1.38$, $p = 0.24$), nor their interaction ($F_{2.74} = 2.63$, $p = 0.08$) had an effect on median RT.

5.3.2 Event-related potentials

Based on known motion-associated ERP components, comparisons between (1) supraliminal coherent versus subliminal random and (2) periliminal coherent versus periliminal random trials were restricted to posterior electrodes. Grand average waveforms and significant spatio-temporal clusters were calculated for each condition and are shown in Figure 5-5 and Figure 5-6 (see 5.2.2.2). Comparing the

ERP waveforms for supraliminal coherent and subliminal random trials (see Figure 5-5), I found a significant cluster (t = 164.89, p = 0.04) between 217 ms and 357 ms after stimulus presentation. During this period, supraliminal trials elicited a more negative response. Looking at the individual electrodes, the effect was restricted to the left hemisphere (averaged over electrodes PO3, PO7 and P3 and non-significant trend in surrounding electrodes). No equivalent significant differences were found in right hemisphere electrodes.

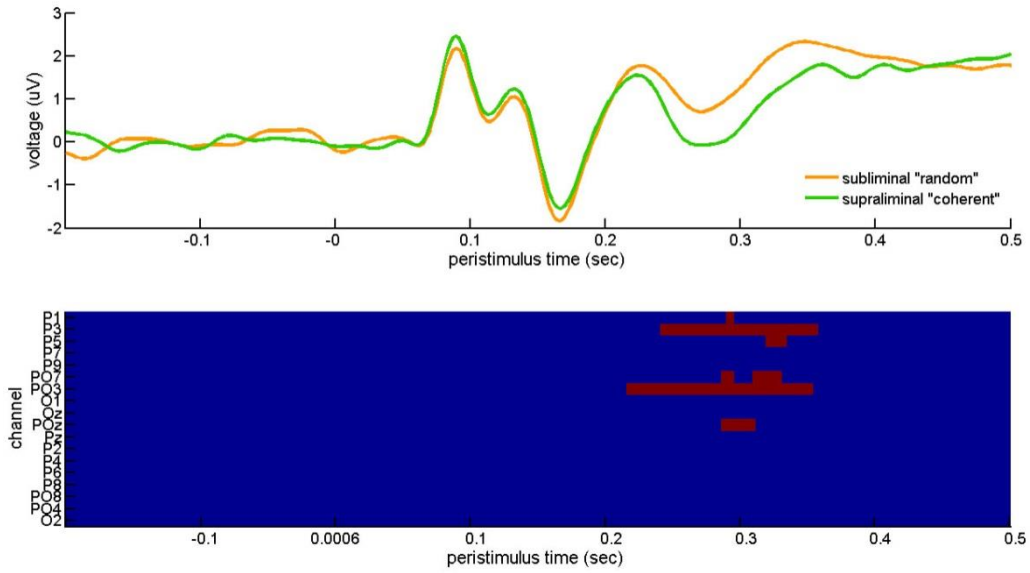

Figure 5-5 Grand averages comparing correct subliminal and supraliminal trials. Top: Lines represent the voltage of the significant cluster (see main text) for subliminal trials that were perceived as random (orange) and supraliminal trials perceived as coherent (green). Bottom: Peristimulus time is plotted against the occipital electrodes included in the analysis; comparisons were restricted to post-stimulus times 100 ms to 500 ms (see 5.2.2.2). Red signals a significant value for the sample by sample paired t-test between conditions (uncorrected).

The same analysis was performed for periliminal trials comparing those perceived as random versus those perceived as coherent. Here, I observed a more negative response (t = 437.97, p = 0.01) for trials that were perceived as coherent compared to

those perceived as random (Figure 5-6) between 207 ms and 500 ms. Note, that these trials were physically identical.

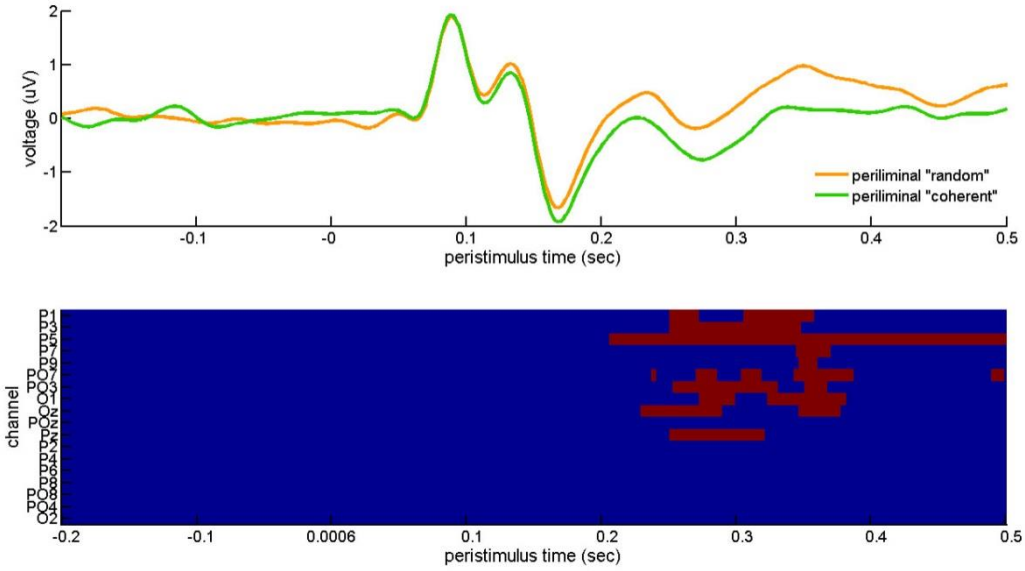

Figure 5-6 Grand averages comparing periliminal trials. Top: Lines represent the voltage of the significant cluster (see main text) for periliminal trials that were perceived as random (orange) and periliminal trials perceived as coherent (green). Bottom: Peristimulus time is plotted against the occipital electrodes included in the analysis; comparisons were restricted to post-stimulus times 100 ms to 500 ms (see 5.2.2.2). Red signals a significant value for the sample by sample paired t-test between conditions (uncorrected).

5.3.3 Time frequency analysis of prestimulus activity

Next, I analysed the prestimulus difference between coherent and random periliminal trials using time frequency analyses looking separately at low (2 – 30 Hz) and high frequencies (25 – 100 Hz). Restricting my analysis to the posterior electrodes, I found one positive cluster for the low frequencies spreading over almost all electrodes except P8 and P9 (although to very different degrees) (t = 383.92, p_{cluster} = 0.04) lasting from -824 ms to -262 ms. The cluster included frequencies between 7 – 15 Hz and was more pronounced in the lateral electrodes. Thus, the cluster indicates

a higher prestimulus alpha activity prior to stimulus onset for periliminal trials that are subsequently perceived as coherent compared to random ones. See Figure 5-7 for an overview of the low frequency results.

For a similar analysis performed on high frequencies, I found three positive clusters that reached significance. The first one (t = 87.78, $p_{cluster}$ = 0.02) was more lateral and most pronounced in the right hemisphere including frequencies between 25 – 55 Hz and lasted from -412 ms to -157 ms. The second one (t = 80.83, $p_{cluster}$ = 0.02) spanned a similar frequency range from 25 – 45 Hz, but was present earlier before stimulus presentation, i.e. from -987 ms to -731 ms. Also, it was more medial and occipital. The third cluster (t = 73.30, $p_{cluster}$ = 0.02) also included medial electrodes and lasted almost up to stimulus onset, i.e. -205 ms to -109 ms, from comprising quite a broad range of frequencies from 35 – 95 Hz. See Figure 5-8 and Figure 5-9 for an overview of the high frequency results.

One cluster comparing prestimulus activity between supraliminal coherent and subliminal random trials was found, i.e. a negative cluster in the low frequency domain (t = -541.48, $p_{cluster}$ = 0.01) spanning frequencies between 13 – 18 Hz and lasting from -859 ms to -109 ms. No significant results were found in the high frequency domain (lowest p = 0.44). This is in line with the assumption that perception in these conditions should be mostly based on post-stimulus instead of on prestimulus factors.

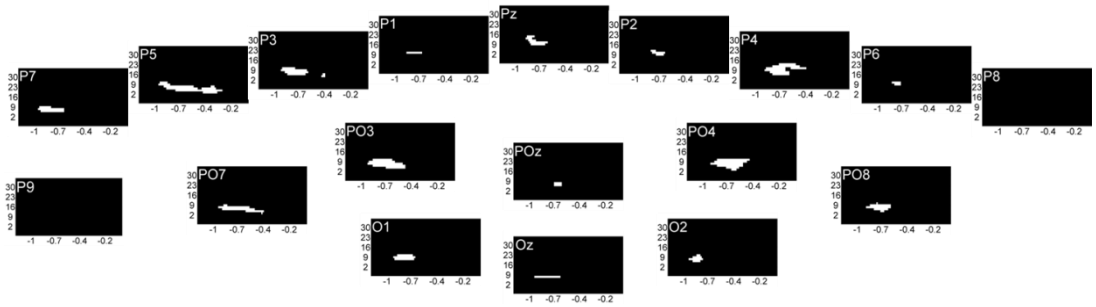

Figure 5-7 Low frequency prestimulus analysis. Higher activity in the alpha range was found for coherent compared to random periliminal trials. One positive cluster was observed including frequencies from 7 – 15 Hz. Single diagrams show separate channels (indicated by the number in the upper left corner) arranged in a similar way to scalp topography. X-axis: prestimulus time (sec), y-axis: frequency (Hz).

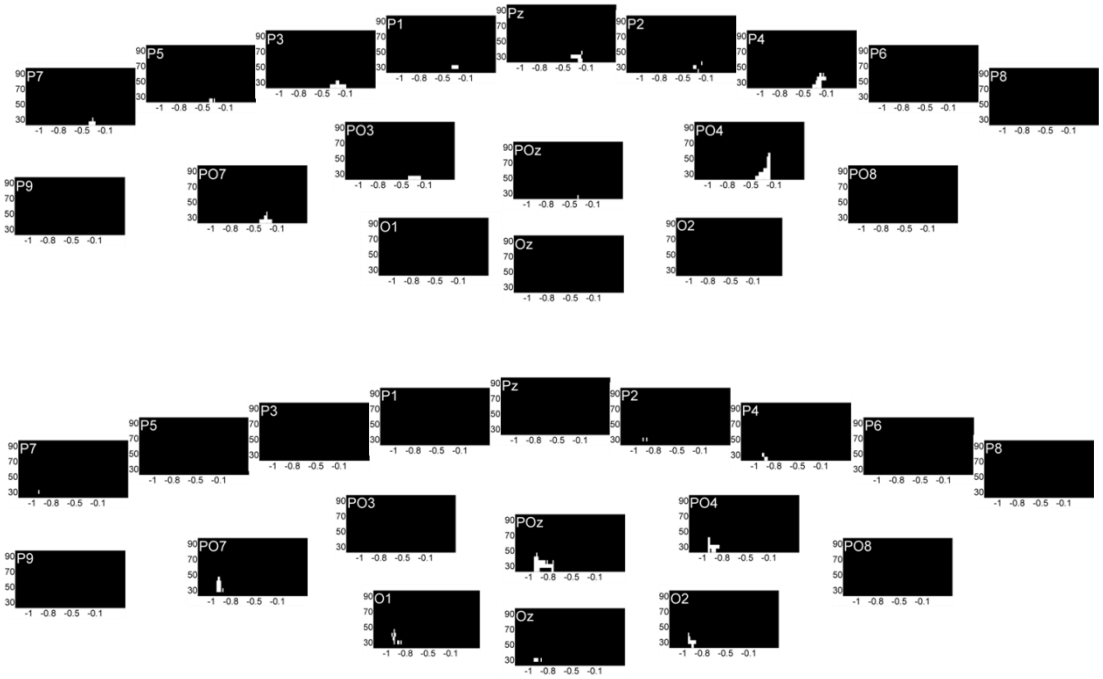

Figure 5-8 High frequency prestimulus analysis. Higher activity in the beta band was found for coherent compared to random periliminal trials. Two positive clusters were observed including frequencies from 25 – 55 Hz. Single diagrams show separate channels (indicated by the number in the upper left corner) arranged in a similar way to scalp topography. X-axis: prestimulus time (sec), y-axis: frequency (Hz).

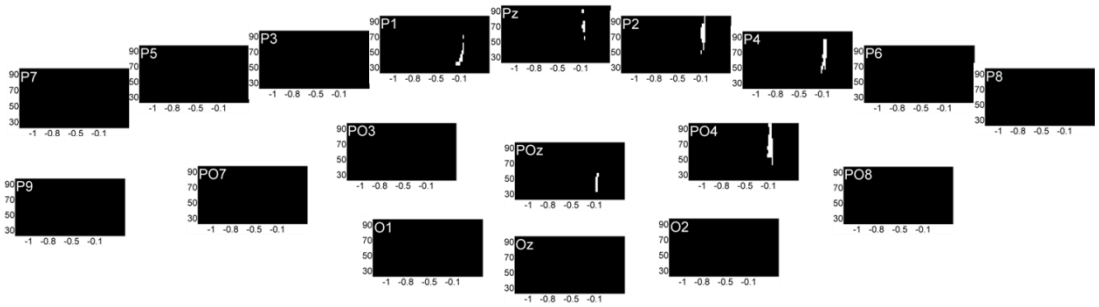

Figure 5-9 High frequency prestimulus analysis. Higher activity in the gamma band was found for coherent compared to random periliminal trials. One positive cluster was observed including frequencies from 35 – 95 Hz. Single diagrams show separate channels (indicated by the number in the upper left corner) arranged in a similar way to scalp topography. X-axis: prestimulus time (sec), y-axis: frequency (Hz).

5.4 Discussion

The aim of this study was to address the question of whether specific signatures of intrinsic oscillatory activity prior to stimulus presentation predicted perceptual outcome in an ambiguous random dot motion coherence task. In order to do so I used three different types of random dot stimuli, i.e. one with high coherence easily detectable by all participants (supraliminal), one with no coherence (subliminal) and one at threshold coherence based on individual detection thresholds (periliminal). I verified my experimental approach by showing that the ERPs elicited by physically different stimuli show the well-documented N2 effect, a motion-induced negativity peaking around 200 ms after stimulus onset found in posterior (occipital) electrode sites (Niedeggen and Wist, 1999; Kuba et al., 2007; Martin et al., 2010). I observed the same pattern comparing periliminal trials that were perceived as coherent to those perceived as random and this corroborated the subjective reports of my participants. My time-frequency analysis of the prestimulus data revealed that both low and high frequencies play a role for the subsequent perceptual outcome of periliminal trials. I found four positive clusters, one in the alpha band, two in the beta band and one in

the gamma band, showing temporally specific higher power in different frequency bands in the prestimulus interval of periliminal trials that were perceived as coherent versus those perceived as random.

5.4.1 ERP results

To validate my experimental approach I first compared the resulting ERPs elicited by trials that were coherent versus those that are random, i.e. physically different (supraliminal versus subliminal) and physically identical trials (periliminal random versus periliminal random). I observed a negativity post-stimulus for coherent trials confirming the importance of the so-called N2 effect, which is considered to be the main motion-specific ERP (Kuba et al., 2007). However, my N2 effect was observed slightly later than the one reported in the literature. This difference might be explained by the different nature of my stimuli compared to those used by Niedeggen and Wist (1999) and others. Ours involved a full stimulus onset, rather than a change from a static to a moving dots display. This difference caused several other early ERP components associated with luminance onsets and therefore might delay motion-related components. Regardless, my ERP results support the interpretation that periliminal coherent and random stimuli received different post-stimulus processing despite having exactly the same visual input.

5.4.2 Alpha band oscillations

My time frequency analysis of prestimulus intervals of periliminal trials revealed one positive cluster in the typical alpha range between 8 and 12 Hz indicating a higher alpha power prior to trials perceived as coherent compared to random ones. The cluster was quite lateralised over the posterior electrodes, was observed early before

stimulus presentation (starting at -824 ms) and did disappear before stimulus onset (at -262 ms). Spontaneous cortical activity is known to express alpha oscillations that are manifested as a visible peak in power spectra of EEG (Berger, 1929; for review, see Klimesch et al., 2007). Moreover, alpha oscillations have been observed during spontaneous conditions and prestimulus intervals. However, the role of alpha power during these intervals remains a controversial topic. The power in this frequency band has been shown to be positively (Linkenkaer-Hansen et al., 2004; Zhang and Ding, 2010) or negatively (Thut et al., 2006; Hanslmayr et al., 2007; Romei et al., 2008) correlated with psychophysical performance[12]. Crucially, the studies reporting a positive correlation between alpha power and stimulus detection used weak stimuli similar to my periliminal trials during which participants indicated the detection of a coherent motion in 50 % of the cases. Thus, my findings are in line with the idea that ongoing alpha oscillations in sensory cortices prior to stimulus presentation can augment the processing of weak sensory stimuli.

5.4.3 Beta band oscillations

In addition to higher power in the alpha band preceding periliminal trials that were perceived as coherent, I also found two clusters in the beta range, typically defined as the range between 12 and 30 Hz. Both clusters were in the higher end of beta, but showed a distinct pattern in terms of distribution and timing. The earlier cluster (from -987 ms to -731 ms) prior to stimulus onset) was more medial and very occipital, whereas the later cluster (from -412 ms to -157 ms) prior to stimulus onset was more lateral and more pronounced in the right hemisphere. My findings are in

[12] A potential explanation for directional differences has been provided by local field potential (LFP) recordings showing that the laminar origin of the oscillatory source might be a distinguishing factor (Bollimunta et al., 2008; Mo et al., 2011).

line with the idea that beta band coherence is important for perceptual decisions in different modalities, e.g. visual (Piantoni et al., 2010), tactile (van Ede et al., 2010) , and sensorimotor (Engel and Fries, 2010). More precisely, my findings provide further evidence for the theory that beta band oscillations are an index of visual perception, showing the most prominent effect between different motion percepts over occipital-parietal areas (Piantoni et al., 2010). In fact, it has been proposed that long-range communication in neural networks of synchronised beta oscillations mediate neural communication and predict perception (Donner et al., 2007; Senkowski et al., 2007; Hipp et al., 2011), i.e. facilitating perception and reporting of a stimulus (Donner and Siegel, 2011). In contrast to the findings of Donner et al. (2007), I find that prestimulus beta band power is predictive of the behaviourally reported outcome. Importantly, I aimed to promote the potential influence of spontaneous fluctuations, whereas Donner at al. tried to maximally control participants' baseline activity. In particular the two studies differed in their interstimulus interval, stimulus presentation time and fixation instructions.

In conclusion, the observed increased beta band activity found here might indicate enhanced top-down attention to coherently perceived trials, thereby leading to an improved detection of coherent motion.

5.4.4 Gamma band oscillations

I also found one significant cluster in the gamma range (from -205 ms to -109 ms), commonly defined as the frequency range from 25 to 100 Hz. This cluster encompassed a large range of frequencies and lasted until stimulus presentation. Previously, high-frequency activity in the human visual motion pathway has been

suggested to play a role for the encoding of visual motion intensity. Using MEG, Siegel et al. (2007) have shown that activity in the gamma range increases with visual motion strength and that this activity originates from visual areas involved in motion processing. Moreover, enhanced gamma band activity in contralateral motor cortex has been shown to predict participant choices in a visual motion detection task several seconds before stimulus presentation (Donner et al., 2009). I underline the importance of prestimulus gamma band activity for perceptual choices in general – and coherent motion detection in particular – by showing that prestimulus gamma in visual brain areas is also predictive of subsequent motion detection.

5.4.5 Conclusion and future direction

Here, I provide evidence that alpha, beta and gamma oscillations indicate a specific state of "readiness" associated with the detection of coherent motion of a participant, and that the increased power of these oscillations prior to stimulus presentation exhibit a specific temporal and spatial pattern predicting motion detection on a single trial basis.

However, using EEG or any other neuroimaging method, no causal inferences about the measured signal can be made. It is possible that the oscillatory patterns are mere epiphenomena of the underlying cognitive mechanisms or driven by an unobserved aspect of brain activity that is driving both my results here and the fluctuations in perception. In humans, transcranial magnetic stimulation (TMS) is the only technique that is used to perturb specific brain regions of healthy individuals. Repetitive TMS (rTMS) can be applied to entrain neural populations with a distinctive frequency (Thut et al., 2011) allowing to test the functional role of

oscillatory synchrony. Thus, rTMS could be used to assess the functional role of the observed oscillations. Targeting visual areas, rTMS should be used at the different observed frequencies – possibly individually defined – to stimulate the cortex prior to stimulus presentation. The same stimuli as described here could be used. Thus, if for example beta band coherence functionally biases perception towards coherently seen stimuli, I would expect to see a higher rate of coherently perceived periliminal trials after rTMS pulses at a frequency between 12 and 30 Hz. This condition could be compared to a control frequency, stimulation at a control site, as well as sham rTMS. Using this strategy, it could be established whether oscillatory activity in specific frequency bands plays a causal role for coherent motion perception.

Chapter 6 The relationship between mind-wandering, creativity and neuronal coupling

6.1 Introduction

The mind tends to wander off, often jumping between thoughts related to past or future experiences. According to some estimates, this introspection occupies up to half of our awake time (McMillan et al., 2013). This tendency to leave the constraints of the perceptual moment and entertain internally generated, spontaneous thoughts is called mind-wandering – sometimes referred to as task-unrelated thought (Levinson et al., 2012) or daydreaming (McMillan et al., 2013). It occurs during all kinds of mental activities (Killingsworth and Gilbert, 2010), preferably during those where vigilance might be low; e.g. driving or reading (Smallwood, McSpadden, and Schooler, 2008; Galéra et al., 2012). Mind-wandering usually leads to compromised performance of an external task (Smallwood et al., 2003; Smallwood, Fishman, et al., 2007). Even though it is not directed toward an external task, it has been considered as a goal-driven process – that might serve multiple adaptive functions, such as future planning and the resolution of prescient issues (Smallwood and Schooler, 2006).

Mind-wandering is both a stable and a transient aspect of an individual: Previous studies have linked the tendency to mind-wander to character traits and states, such as the disengagement in a sustained attention task (Smallwood et al., 2004). The propensity to mind-wander seems to be a stable cognitive characteristic that predicts performance and difficulties in daily life and in the laboratory (McVay et al., 2009),

happiness (Killingsworth and Gilbert, 2010) and the scores for the Big Five personality traits (Zhiyan and Singer, 1997), such as ‘openness’.

One particular trait of interest, with regard to mind-wandering, is creativity. Several lines of research have linked mind-wandering to enhanced creativity, especially for problems that have been encountered previously. For an overview, see Baird et al. (2012). Baird et al. (2012) were the first to study mind-wandering during a so-called incubation task. An incubation task involves a period during which the individual refrains from task-related thought in the process of problem solving and either engages in an alternative task, or simply rests. Even though there are plenty of anecdotes about Eureka moments occurring during unrelated trains of thought, the scientific investigation of this – potentially crucial – source of inspiration is difficult. Using an incubation task and a validated creativity task, the Unusual Use Task (UUT), Baird et al. (2012) showed that creative problem solving for previously seen items increased, if participants performed an undemanding external task during the incubation period. These increases were observed in comparison to individuals who did not perform any task and those who performed a more demanding task. The group performing the undemanding task also showed higher scores of mind-wandering. Thus, engaging in a simple external task – that allows the mind to wander – facilitates creative problem solving.

The most challenging issue, when studying mind-wandering, is how to measure its prevalence. In addition to the usual confounds related to participants’ reports (e.g. questionnaires), mind-wandering often occurs without awareness (Schooler et al., 2011) and is interrupted by verbal sampling. In reading studies, eye movements,

fixations and blinks are a good indicator of attentional lapses that are characteristic of a wandering mind (Smallwood, McSpadden, and Schooler, 2008; Smilek et al., 2010). However, they might be confused with a lack in understanding that requires the re-reading of a certain part of the text. In other behavioural and neuroimaging studies, two methods have been used to quantify the extent an individual is, or has been, engaged in mind-wandering: either a post-task enquiry about the extent of mind-wandering or experience sampling is used. The latter method is applied at different times throughout the task and can either measure fluctuations of mind-wandering over a longer period or separate samples can be combined to provide a more robust estimate. Because mind-wandering can occur with and without awareness, experience sampling can include a second question about the awareness of thought content; i.e. asking whether the individual was aware of whether the mind was wandering.

The neural substrates of mind-wandering are only beginning to be unravelled. Mind-wandering is associated with activity in the DMN and an individual's report of her tendency to mind-wander is correlated with activity in the same regions (Mason et al. 2007). Thus, thought sampling outside of the scanner and tasks that had been practiced to different levels can be used to investigate the relationship between brain activity, mind-wandering and task novelty. Christoff et al. (2009) extended these findings, proving the first study that employs experience sampling to measure mind-wandering during fMRI. Indeed, it seems as if mind-wandering is a unique state, during which regions of the DMN and task-active or executive network regions are recruited. Furthermore, the activity in the two complementary networks seems to be strongest when individuals are not aware of their wandering minds. The construct of

mind-wandering as a unique mental state – that may allow otherwise divergent brain networks to cooperate – has been further supported more recently: the generation and maintenance of an internal thought can activate DMN regions, as well as a frontal-parietal control network helping to sustain and buffer internal trains of thoughts against disruption from outside (Smallwood et al., 2012). In conclusion, the distinction between DMN and task, attentional or executive networks is not an "either or" one, where one is silent when the other one shows increased activity. It rather seems to be the case that the different networks are capable of integrating seemingly disparate reports of their role in goal-directed behaviour and mind-wandering (Esterman et al., 2012).

In this study, I was interested in establishing the missing neurobiological link between creativity and mind-wandering. My hypothesis was that there is phenotypic variation over individuals that predisposes them to being more or less creative – and that this variation is reflected in the coupling between intrinsic brain networks associated with introspection; i.e. the DMN, and task related networks respectively. I therefore measured this coupling in terms of effective connectivity among distributed cortical regions using fMRI. I hypothesised that subject to subject variations in the coupling between the two networks would be correlated either with (i) a creative trait, (ii) state dependent measures of creative problem solving as assessed before and after scanning or (iii) both. To distinguish between the phenotypic trait of creativity and the states of creative thinking, I measured creativity based upon the participants' personal history and performance on the standard UUT task during scanning. If there is a creative phenotype, I hoped to see significant correlations between subject specific creativity scores and coupling between their brain networks. Conversely, if

creativity is a state dependent product of distributed processing, I hoped to see correlations between measures specific to the resting state (such as performance on a creativity task and mind-wandering during the resting state) and connectivity. Finally, I hypothesised that the creative trait will be expressed in terms of state dependent measures either at the behavioural or the physiological level. In other words, subjects who are more creative will show greater mind-wandering and creativity performance, with or without physiological correlates. In summary, my aim was to establish a link between creativity, creative task performance and mind-wandering at the behavioural level, and to validate these constructs by showing that one or more of these behavioural phenotypes is associated with physiological coupling between the DMN and a task-active network.

6.2 Materials and methods

My paradigm was designed to measure the interactions or coupling between the default mode and task positive or executive networks in the brain. I therefore used fMRI to identify notes belonging to the DMN that were activated by the incubation task, i.e. an incidental target detection task. Furthermore, I was interested in the state dependent effects of this coupling and creativity. I therefore assessed creativity in terms of the improvement in the UUT before and after the incubation task. Scanning was performed during the incubation period involving experience sampling and subsequently to the UUT during a block task contrasting the same incidental target detection task as used during incubation. The block task was used to identify task-responsive regions. The two scans are hereafter referred to as incubation task (i.e.

between the two UUT sessions) and block task (i.e. after the UUT) respectively (see Figure 6-1).

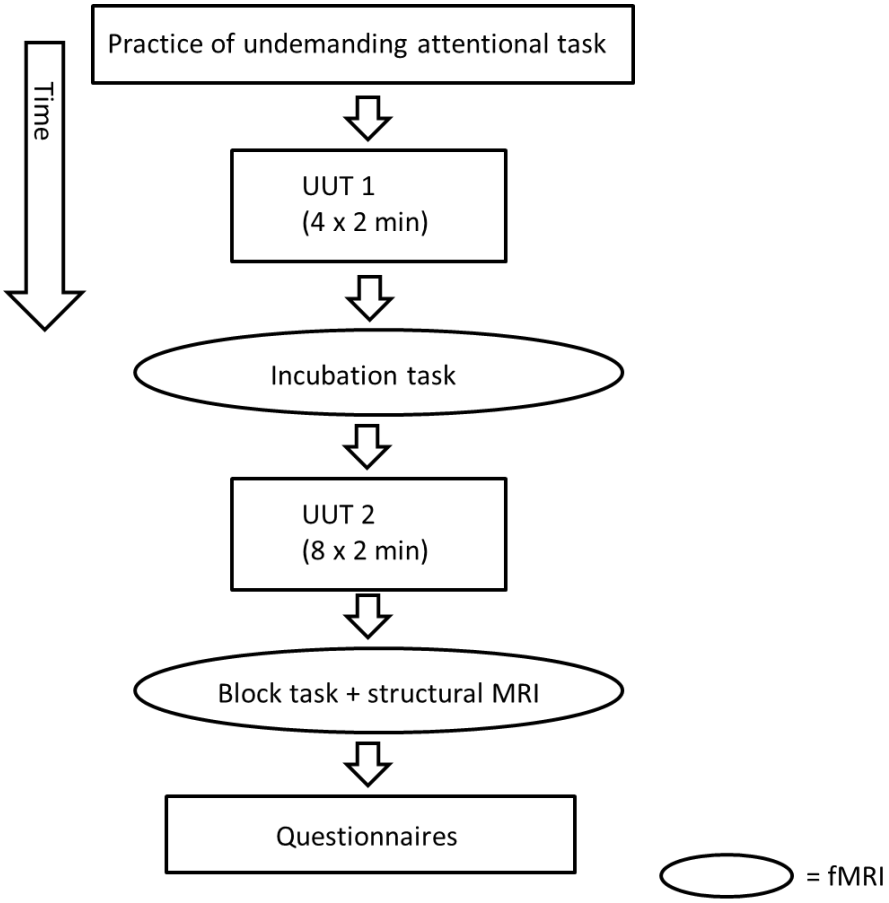

Figure 6-1 Experimental paradigm. Participants first practiced the attention task used for the incubation and the block task. After the first UUT, they performed the incubation task in the scanner and after the second UUT they performed the block task in the scanner. The whole experiment lasted about 2 hours.

6.2.1 Participants

26 right-handed healthy volunteers (9 male, 20-46 of age range, mean age 25.0 years) with normal or corrected to normal vision gave written informed consent to participate in the study – comprising a combined scanning and behavioural session lasting for approximately two hours. One participant was excluded from the analysis

because she fell asleep during scanning and another participant was excluded because she had a Beck Depression Inventory (BDI) (Beck et al., 1996) score of 41 and on subsequent questioning confirmed that she had been diagnosed as being clinically depressed. The study was approved by the local ethics committee.

6.2.2 Stimuli and task design

We employed an incubation paradigm, during which participants were scanned while performing an undemanding attentional task that was used to trigger mind-wandering (Baird et al., 2012). During this task numbers between one and nine were presented centrally for 1 s on the screen followed by a 1 s fixation / response period. 8 % of the trials were targets; i.e. coloured numbers, whose occurrence had to be indicated via a button press. Also, experience sampling reports of mind-wandering were collected during fMRI scanning using thought probes (Christoff et al., 2009). These provided subjective reports of mind wandering. Each probe asked the participants two questions about their mental state directly preceding the probe. The first question asked to what extend the participants' attention had been focused on the task. The idea of mind-wandering as "any thoughts that are unrelated to the task" was explained prior to the experiment; examples of thoughts that are related to the past, the future, worries or other people were also provided. The second question asked to what extend the participants had been aware of where their attention was focused. "No awareness" was defined as not recognising that mind-wandering had occurred until the moment the probe was presented (Schooler et al., 2004; Smallwood, McSpadden, et al., 2007; Smallwood, McSpadden, Luus, et al., 2008). Examples for this kind of situation were given prior to a training session of the task. Participants

answered using a 7 - point Likert scale, ranging from “completely on task” to “completely off task” for the first question, and from “completely aware” to “completely unaware” for the second question. The scale direction was counterbalanced across participants. The 26 targets – comprising 8 % of all trials – and the 12 probes occurred pseudo-randomly following a Poisson distribution at moments the participants could not predict.

The UUT, a widely used measure of convergent thinking (Guilford, 1967), was used to test creativity immediately before and after scanning during the incubation period. I selected the UUT because it has been shown to yield robust and consistent incubation effects (Ellwood et al., 2009; Sio and Ormerod, 2009), compared to divergent-thinking tasks, like the Remote Association Task that are more prone to empirical inconsistencies (Vul and Pashler, 2007). During the UUT I asked participants to generate as many unusual uses as possible for common objects in a set amount of time: i.e. 2 min per object. Each object was presented together with its normal use, e.g. "brick, used for building". Handwritten answers were provided on separate pieces of paper for each item.

6.2.3 Experimental procedure

To address potential correlations between neuronal connectivity and behavioural measures related to mind-wandering and creativity, participants underwent a single two hour experiment combining behavioural tasks and neuroimaging using fMRI. All participants were naïve to the task. First, they received the above explanation about the attentional task used during the incubation period and for the block task. The test run (12 min) included the probes. All participants were offered an additional test run;

all confirmed that they had received sufficient practice and understood the task after the first round. Subsequently, they received a short explanation of the UUT – by means of an example object; e.g., "newspaper, used for reading" and had the chance to ask questions about the task. Then, every participant received the same four words one by one (i.e. brick, key, pencil, bedsheet) finishing the first round of the UUT prior to scanning. After completion of the incubation task in the scanner, participants completed eight words of the UUT outside of the scanner; including the objects they already saw prior to scanning, as well as four new ones (i.e. brick, chair, key, shoe, pencil, safety pin, bedsheet, button). Then participants were scanned again performing the block task that contrasted the same attentional task as used during the incubation task with fixation periods. Finally, a series of questionnaires were completed by each participant: the Beck Depression Inventory (BDI) (Beck et al., 1996), the Creative Achievement Questionnaire (Carson et al., 2005), and the Daydreaming Frequency subscale of the Imaginal Process Inventory (IPI), which accesses an individual's propensity to mind-wander (Singer and Antrobus, 1972).

6.2.4 Behavioural analysis

6.2.4.1 UUT

To access creative fluency, responses to the UUT were rated by two independent raters blind to condition as either valid or invalid. The inter-rater classification of non-redundant responses, rated according to the manual of the UUT (Guilford et al., 1960) was highly reliable (alpha = .95). For each individual, the scores of the two raters were averaged. Based on the scores for each item the total improvement was calculated per subject for old words. For new words, total scores were counted,

because – by definition – improvement scores cannot be calculated for a single sample.

6.2.4.2 Thought probes

The scores of the thought probes during the incubation task, ranging between one and seven, were averaged for each of the two questions resulting in a "mean mind-wandering" and a "mean awareness" score per subject. Standard deviations of these two scores were also calculated.

6.2.4.3 Target detection

For the targets of the incubation task and the block performance (% correct) and mean reaction time were calculated.

6.2.5 fMRI data acquisition

A 3 T Trio MRI Scanner (Siemens Medical Systems, Erlangen, Germany) with a 32 channel head coil was used to acquire functional data with a standard echo planar imaging (EPI) sequence (matrix size 64 × 74; field of view 192 × 192 mm; in plane resolution 3 × 3 × 2 mm; 48 slices in ascending acquisition order; echo time 30 ms; acquisition time per slice 70 ms; TR 3.36 s). The incubation task comprised 226 volumes, and the block task comprised 126 volumes. After both scanning sessions, B0 field maps were acquired to correct for geometric distortions in the EPI images. Also a structural T1-weighted scan was acquired (matrix size 256 × 240; field of view 256 × 240 mm; in-plane resolution 1 mm × 1 mm; 176 sagittal slices of thickness 1 mm; echo time 2.48 ms; acquisition time per slice 7.92 ms). During scanning, respiration volume and cardiac pulse were measured using a breathing belt placed around the participants' waist and an MRI compatible pulse oximeter attached

to one of the fingers. These data, together with scanner slice synchronization pulses, were sampled using Spike2 (Cambridge Electronic Design Limited, Cambridge, UK) and used for physiological noise correction.

6.2.6 fMRI data analysis

6.2.6.1 Pre-processing

Functional data were analysed using SPM12 beta (http://www.fil.ion.ucl.ac.uk/spm/software/spm12/) and DCM12 was used for dynamic causal modelling of effective connectivity. To allow for T1 equilibration, the first five images of each scanning session were discarded. Pre-processing of the data involved mean bias correction, realignment of each volume to the first volume of each run, coregistration of the functional data to the structural data, normalisation to the MNI template brain and smoothing by an 8 mm Gaussian kernel.

6.2.6.2 Block task

The block task data were filtered with a standard 128-s cut-off. Regressors modelling the task blocks were formed by convolving boxcar functions encoding the condition with a canonical hemodynamic response function. Movement parameters were included as confounds in the first level analysis for each participant. No global signal regression was performed. A standard contrast of the first level parameter estimates was used to perform a random effects analysis over participants in the usual way. This involved estimating (the contrast of) parameters encoding the effect of interest using a standard linear convolution model at the first (within-subject) level and then passing the resulting contrast images to a one sample *t*-test at the second (between-subject) level. The resulting statistical parametric map (SPM) was used to test for

task-active regions. The anatomy toolbox (Eickhoff et al., 2005) was used to anatomically designate activated areas.

6.2.6.3 Incubation task

The data of the incubation task were filtered with a 256-s cut-off, high-pass filter to remove low-frequency drifts, while preserving as many of the spontaneous fMRI fluctuations as possible (Birn, 2007). Because I wanted to model data showing physiologically relevant resting-state (i.e. low frequency) dynamics from regions of the DMN, I used a similar approach to that described in Fransson (2005). Therefore, I modelled fluctuating haemodynamic signals with a discrete cosine basis set consisting of 189 frequencies characteristic of resting-state dynamics (0.0078 – 0.1 Hz) (Biswal et al., 1995; Fox and Raichle, 2007; Deco et al., 2011), using a general linear model. Physiological data (respiration and heart beat) were modelled using an in-house developed MATLAB toolbox (Hutton et al., 2011) based on RETROICOR (Glover et al., 2000). This resulted in a total of 17 regressors. The resulting regressors were included as confounds in the first level analysis for each participant. Movement parameters were also included as confounds. Global signal regression was performed.

6.2.6.4 ROIs

Task active regions were extracted from the block task data: i.e. the peak voxels at $p < 0.05$, FWE corrected, of clusters larger than 10 voxels were used as ROIs for the DCM analysis. Participant-specific coordinates of these ROIs were defined as the closest maximum to the group peak voxels within the specific anatomical region, as defined by the SPM Anatomy toolbox (Eickhoff et al., 2005).

DMN ROIs were identified using a contrast across all discrete cosine basis set regressors and the resulting SPM was masked by the anatomical masks of four classical DMN regions; i.e. ventromedial prefrontal cortex (vmPFC), posterior cingulate cortex (PCC), as well as left and right inferior parietal lobule (Mantini and Vanduffel, 2013). Individual coordinates were defined as the peak responses in the respective brain area (again using the SPM Anatomy toolbox, as well as the WFU Pickatlas (Maldjian et al., 2003) for those regions not labelled by the SPM Anatomy toolbox). Having identified the nodes or ROIs corresponding to the task-active network and the DMN, regional summaries of activity were used for subsequent dynamic causal modelling of effective connectivity within and between the two networks.

6.2.6.5 Dynamic causal modelling

DCM models neuronal dynamics in terms of directed and reciprocal influences among brain regions. Stochastic DCM allows one to model spontaneous or endogenous (non-controlled) activity. It does not require any input usually associated with experimental manipulation. Six subject-specific ROIs defined by the block task and the incubation task data using discrete cosine basis set functions were used as the nodes for a fully connected model for each participant; i.e. with bidirectional extrinsic connections between all regions. Regional activity in each ROI was summarised with its principal eigenvariate (adjusted for nuisance variables) based on voxels within 8 mm of subject-specific peaks. The models were fitted to the time series from the incubation task using generalised (Bayesian) filtering (Li et al., 2011). The resulting estimates of effective connectivity between task active and default nodes were then used to summarise neuronal coupling on a per subject basis.

At the between subject level, I used standard multiple linear regression methods to characterise any differences in coupling between the DMN and the task positive regions and the behavioural (state or trait) measures; i.e. mean mind wandering and mean awareness during the incubation task, UUT improvement for old words, UUT scores for new words, frequency of daydreaming (questionnaire), and creative achievement (questionnaire). These associations were tested using analysis of covariance (ANCOVA), with effective connectivity estimates as independent or explanatory variables and behavioural measures as dependent or response variables.

6.3 Results

6.3.1 Behavioural results

Target detection during the block task was 100 % for all participants; mean reaction time was 0.80 s (range 0.75 – 0.84 s). Mean target detection during the incubation task was 98.4 %; mean reaction time was 0.82 s (range 0.78 – 0.85 s). These results indicate that all participants paid attention during task performance.

Mean mind-wandering score during the incubation task was 4.14 (range 1.75 – 5.75), while mean awareness was 3.69 (range 1.75 – 6.33). These results indicate that participants showed substantial inter-individual variation – as required for the tests of association between the various measures at the between subject level (see Figure 6-2).

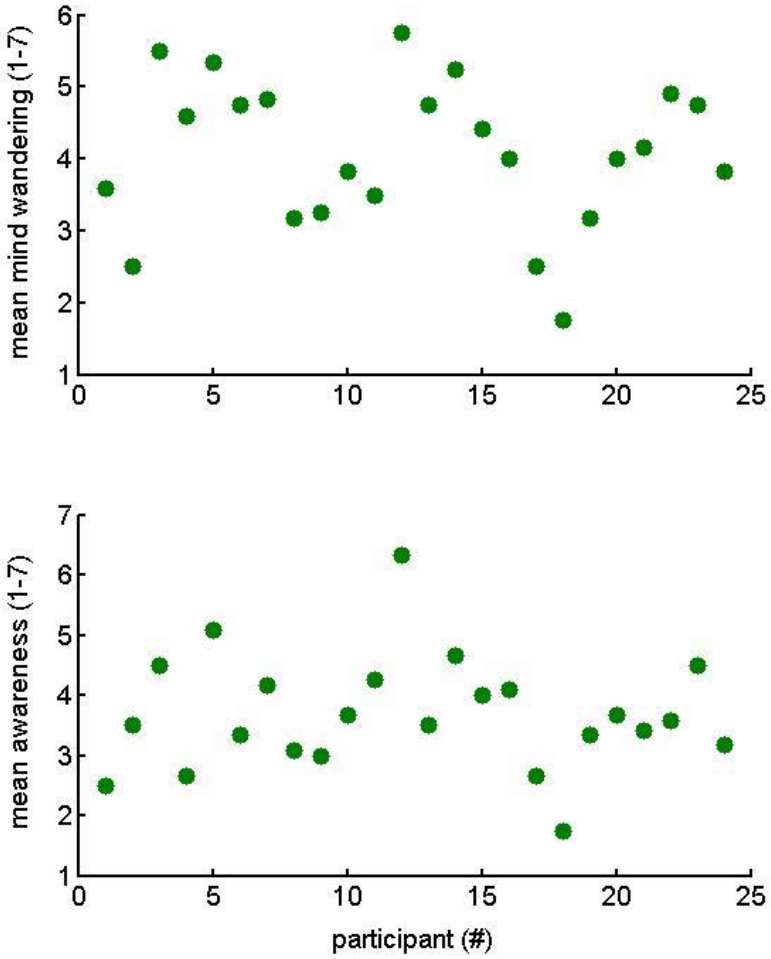

Figure 6-2 Inter-participant differences in mind-wandering and awareness. Participants mind-wandering and awareness scores (averaged over the 12 probes) spanned the whole response space indicating substantial variability over participants.

The following correlations between behavioural measures and questionnaires were observed. Mean mind-wandering and mean awareness (and their standard deviation over the 12 probes) during the incubation task were positively correlated ($r = 0.75$, $p = 0.00002$). The more people mind wandered, the more they were aware of the fact that their thoughts were wandering off. In addition, mean mind-wandering scores

were positively correlated with the frequency of daydreaming in general as indicated by the frequency of daydreaming subscale of the IPI (Singer and Antrobus, 1972) (r = 0.45, p = 0.03). This provides a validation of experience sampling in this paradigm and suggests that participants who daydream more in general, also showed more mind-wandering during the experiment.

The average total improvement for old UUT words was negatively correlated with mean awareness during the incubation task (r = -0.40, p = 0.05) indicating that participants who were *less* aware of whether they were mind-wandering or not, showed significantly higher improvements for the creativity task.

Furthermore, I found a positive correlation between scores of the creative achievement questionnaire and the general frequency of daydreaming (questionnaire) (r = 0.42, p = 0.05) providing evidence for the notion that individuals who daydream more are more creative.

For the scores of the new UUT words, I observed a positive correlation with the amount of daydreaming in general (questionnaire). However, this correlation did not reach significance, but showed a trend (r = 0.38, p = 0.08), indicating that individuals who are more creative in a more general sense (i.e. for new problems) might daydream more in general. In summary, there were strong correlations among state dependent measures of mind-wandering and creativity and between trait measures – with trend correlations between state and trait measures. The key question now is whether these measures could be predicted by neuronal coupling as measured during incubation.

6.3.2 Imaging results

6.3.2.1 Task-active regions

After pre-processing, I used the block task to define regions that are responsive to the undemanding attentional task used for the incubation period. I found bilateral fusiform gyrus (left: -44, -68, -18; right: 40, -66, -16), both at $p < 0.05$, FWE corrected, with the constraint of clusters larger than 10 voxels (see Figure 6-3). Individual coordinates were determined using the methodology outlined in 6.2.6.4 and are shown in Table 6-1.

Table 6-1 Individual coordinates of task related activation in the fusiform gyrus. Individual coordinates were extracted from the block task data and defined as the closest maximum to the group peak voxels within the specific anatomical region.

	Left fusiform			Right fusiform		
Participant	**X**	**Y**	**Z**	**X**	**Y**	**Z**
Group	-44	-68	-18	40	-66	-16
1	-44	-60	-16	46	-64	-18
2	-42	-64	-18	44	-56	-16
3	-38	-60	-16	38	-60	-16
4	-44	-60	-16	38	-80	-18
5	-38	-56	-20	34	-64	-12
6	-52	-50	-24	40	-60	-16
7	-42	-56	-16	44	-60	-16
8	-32	-70	-14	42	-60	-16
9	-34	-58	-12	32	-46	-18
10	-36	-46	-12	42	-56	-24
11	-34	-62	-12	36	-58	-12
12	-38	-60	-16	38	-68	-18
13	-50	-60	-22	34	-62	-12
14	-38	-58	-16	42	-52	-18
15	-42	-76	-20	34	-70	-20
16	-42	-54	-18	40	-52	-16
19	-46	-64	-18	48	-64	-18
20	-38	-60	-16	40	-64	-20

21	-46	-62	-18	50	-60	-16
22	-34	-62	-12	40	-64	-18
23	-42	-60	-16	38	-68	-18
24	-40	-60	-16	40	-62	-18
25	-38	-58	-20	46	-56	-20
26	-44	-62	-18	36	-58	-12

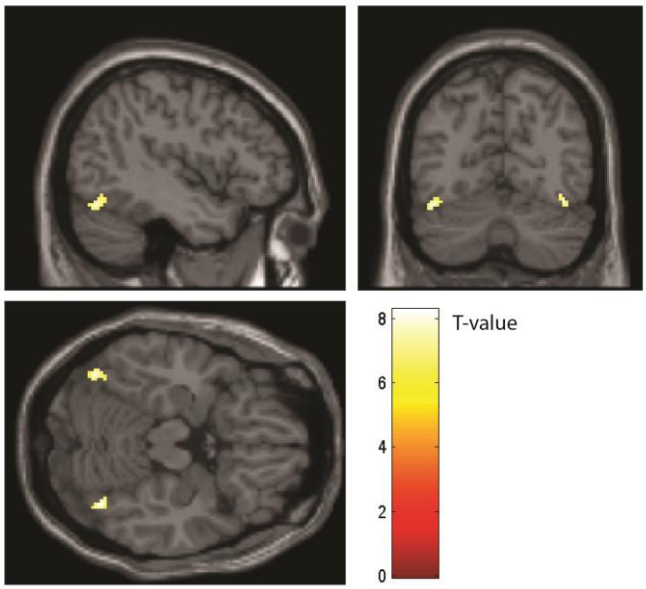

Figure 6-3 Task activations in the fusiform gyrus. The attention task activated bilateral fusiform gyrus, $p < 0.05$ FWE corrected.

6.3.2.2 DMN regions

Having identified bilateral fusiform gyrus as the key regions that show task-related activation during the attention task (using the data of the block task), I next identified the ROIs of the DMN, using the timeseries from the incubation task – model with a discrete cosine basis set.

6.3.2.3 Stochastic DCM

My subsequent tests for correlations between effective connectivity and behavioural measures used (within subject) effective connectivity estimates from stochastic DCM

in a (between subject) classical ANCOVA. My hypotheses were not about the existence of connections, but whether coupling between the task positive and DMN nodes could predict: (i) measures of mind-wandering during the incubation task, (ii) improvements during the UUT or (iii) creative achievement. Therefore, I calculated three different ANCOVAs for each behavioural measure, examined whether these were significant and identified the connections that were driving any significant associations.

Only the CAQ gave significant results (p = 0.013, uncorrected), which were mainly driven by two connections: from PCC to left fusiform gyrus (68.5) and from PCC to right fusiform gyrus (-60.5). Two other connections showed a moderate contribution: IPL to left fusiform gyrus (45.3) and fusiform right to IPL right (-37.3). See Figure 6-4 for the correlation and contribution of the different connections. These contributions correspond to the regression coefficients implicit in ANCOVA, when regressing creativity on effective connectivity. In brief, higher connection strength from DMN nodes to the left fusiform node relative to the right fusiform node indicates that participants are likely to report more creative traits.

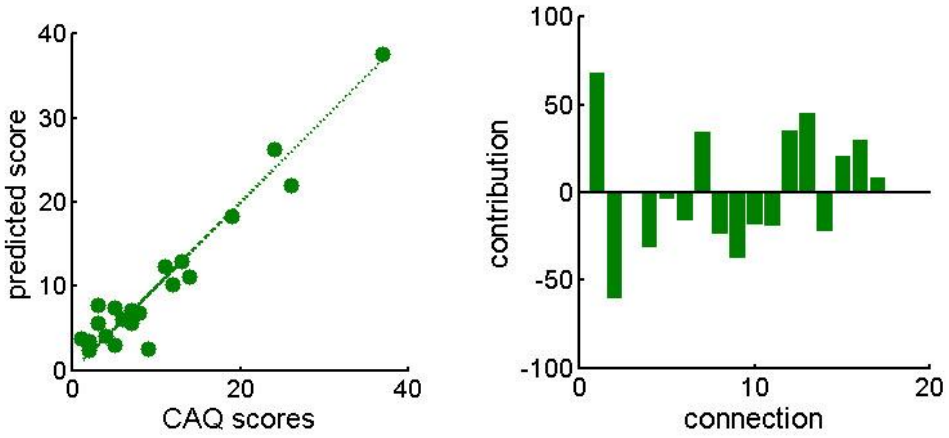

Figure 6-4 Correlation between creativity and brain connectivity. Left: CAQ scores were correlated with the predictive score of the coupling. Right: The sixteen connections between regions of the DMN and the task-positive network showed a different contribution to the correlation shown on the left side. The first two are from PCC to left fusiform and from PCC to right fusiform respectively.

6.4 Discussion

I was interested in establishing the missing neurobiological link between creativity, mind-wandering and neural coupling. Therefore, I tested the hypothesis of whether there is a phenotypic variation in creativity traits over individuals. I hoped that this variation would be reflected in the coupling between intrinsic brain networks associated with introspection, i.e. the DMN, and task-related processing respectively.

I found that behavioural measures of state and trait creativity indeed were correlated with the tendency to mind-wander or the awareness about it respectively during an incubation task, as well as the general tendency to mind-wander or daydream in life. On the neuronal level, I found that the coupling between the two networks was correlated with the creative trait as measured by creative achievement, but not with state-dependent measures of creative problem. Thus, my results indicate that there is

a creative phenotype that is reflected in the coupling between the two brain networks. Furthermore, this creative trait was expressed in terms of another behavioural trait, namely the tendency to mind-wander during daily life. Thus, I established a link between creativity, creative task performance and mind-wandering at the behavioural and neurophysiological level.

6.4.1 Behavioural results

I used an undemanding incubation task and a standard behavioural creative problem solving task to address convergent thinking and measured state creativity in two ways. First, the ability to find creative solutions to "old problems"; i.e. objects that had been encountered previously, and second the ability to find creative solutions to "new problems"; i.e. objects that participants did not see before. With regard to the measurement of mind-wandering during the incubation task, I used experience sampling asking participants at random time-points to measure whether their minds were wandering and their awareness of this wandering. This approach allows a much more accurate picture of the mental state of participants than a simple questionnaire administered after the task (Christoff et al., 2009; McVay et al., 2009). I was particularly interested in the distinction between mind-wandering and the awareness of mind wandering. Earlier studies have shown a clear dissociation between those two processes (Smallwood, McSpadden, et al., 2007; Christoff et al., 2009; Schooler et al., 2011).

Another distinction provided by my ratings is between task-independent and stimulus-independent thoughts. According to this distinction, mind-wandering is only present if thoughts are independent of stimulus *and* task; i.e., a mere distraction

of the task at hand triggered by a sudden noise or some background drumming does not fall into the category of mind-wandering (Stawarczyk et al., 2011). I confirmed and extended the findings of Baird et al. (2012), who showed that individuals who mind wandered more during an incubation task showed improved creative problem solving for revisited items, but not for new items. However, they only used a post-task questionnaire to assess mind-wandering during the task. Such a measure runs into the risk of ignoring large parts of the data, because of the recency effect (Ebbinghaus, 1913; Baddeley and Hitch, 1993). Furthermore, the authors did not distinguish between mind-wandering and awareness. Here, I found that the performance on old problems of the UUT task was negatively correlated with mean awareness about where an individual's attention was focused. In other words, the less we are aware of where our mind is, the better our results on a state dependant measure of creativity. Also, I observed a positive correlation between the mean amount of mind-wandering and the mean awareness scores indicating that individuals who mind-wander more are – on average – more aware about whether their mind is occupied with things outside of the external world. However, the strongest neural activations are observed for periods of mind-wandering without awareness (Christoff et al., 2009).

My trait measurement of creativity was correlated with my trait measurement of mind-wandering; providing evidence for the idea that individuals who mind-wander more are more creative in general. Crucially, my trait measurement of creativity neither tested for creative performance in a practical and direct way like the UUT or non-verbal tasks included in the Torrance Test of Creative Thinking (TTCT) (Torrance, 1962), nor did it ask participants about how creative they think they might

be. Both accounts suffer from uncertainties that can either distort performance and/or accuracy of how creative somebody truly is. In order to get a valid measurement of creativity I addressed creative achievement over the life span including areas from scientific success to cooking and painting. In summary, my results are in line with the previous literature establishing a link between mind-wandering and creativity, where individuals who are more creative tend to mind-wander more often (Dijksterhuis and Meurs, 2006; White and Shah, 2006; Sio and Ormerod, 2009; Baird et al., 2012).

6.4.2 Imaging results

On the neuronal level, I have established a link between the coupling of two intrinsic networks that are usually associated with very diverse contexts; namely, mind-wandering and task-execution. I show that the coupling between these two networks is increased (in the left hemisphere) for participants who show a more creative phenotype. This correlation was mostly driven by two connections, namely from PCC to bilateral fusiform gyrus.

Until now, the neural basis of mind-wandering has only been investigated in very few studies. Because of the difficulty with measuring mind-wandering and its proneness to individual judgements, recency effects and other confounds, it is difficult to compare even these studies. One finding, that has been observed with consistency is the involvement of the DMN network in the process of mind-wandering (Christoff et al., 2009; Stawarczyk et al., 2011; Esterman et al., 2012), as well a task-activated network; e.g., the dorsal attention network (DAN) (Esterman et al., 2012) or the executive network (Christoff et al., 2009). However, all previous

studies were based on activation studies and did not investigate the potential role of coupling between these two networks. I provide the first evidence showing that coupling between these two networks is indeed linked to a creative phenotype. The PCC is one of the core regions of the DMN (Shulman et al., 1997; Raichle et al., 2001) and might form – together with the subregion of PFC involved in the DMN – the core of the DMN (Raichle et al., 2001). Furthermore, this region has been described as a tonically active brain region that is constantly gathering information about the external world – and possibly the internal one (Raichle et al., 2001). In light of these considerations, the PCC might also be the most strongly involved in the modulation of purely task-active regions, like the fusiform gyrus. In short, it might provide a bridge between task positive and task negative during mind-wandering.

6.4.3 Limitations

There are a few limitations of this study that should be mentioned. First, the individual scores provided by the experience sampling are subject to several sources of uncertainty. Even though, all participants received a detailed explanation of the task and performed a training session, I cannot address individual decision criteria determining their ratings. For instance, two individuals might be in the exact same state in terms of the amount they are mind-wandering (at least theoretically), but might rate this condition to a different degree. These decision criteria depend on both state and trait variables. As mentioned earlier, mind-wandering is not an on-off condition, which can also be confused by purely task-unrelated thoughts that are related to other external stimuli, like scanner noise or an itching nose (Stawarczyk et al., 2011).

Second, creative achievement depends on age. A first year student simply does not have that many years to publish papers, win acting prizes or write a book. Even in a relatively restricted age range sample like mine, this factor needs to be taken into account.

Third, I used the mean of the twelve experience sampling ratings of mind-wandering and awareness in order to get a robust summary measure. However, by doing so I lost the fluctuations in mind-wandering throughout the incubation period. Previous work has pointed out that within subject differences in the amount of mind-wandering and awareness lead to differences in the neural activity of brain regions of the DMN and task-executive regions (Christoff et al., 2009). Of course, taking such an approach leaves the arbitrary decision about which time interval to look at; in other words, how early or late prior to the experience sampling imaging data should be included in the analysis.

Last, the definition of ROIs always brings certain limitations, because they depend on several variable settings that add noise. The location and size of each ROI were chosen carefully in this study, but choices often depend on empirical values that might not be appropriate for the study at hand.

6.5 Conclusion

In conclusion, I have established a link between creativity, creative task performance and mind-wandering on the behavioural level and validated these constructs by showing that creativity is associated with physiological coupling between the DMN and a task-active network stimulated during a standard incubation task.

Chapter 7 General discussion

The underlying neurobiology of spatially and temporally correlated patterns of ongoing brain activity is still poorly understood (Schölvinck et al., 2013). Also, its function and purpose have been much debated during the last years (Boly et al., 2009; Honey et al., 2009; Sadaghiani, Hesselmann, et al., 2010; Raichle, 2011). The work presented here is one line of research targeted to understand the interaction between spontaneous brain activity and behaviour, both on fast and a longer temporal scale. These include prestimulus intervals both in the healthy and the lesioned brain, as well as learning experiences and behavioural phenotypes. In sum, this work is one of the numerous small steps on the way to gain deeper insight into the functioning of the brain as a whole including its extensive network of functional connections.

7.1 Overview of findings

Using functional imaging to measure spontaneous fluctuations in resting state activity, I first demonstrated that long-lasting changes in neuronal coupling can be accompanied by changes in resting state activity. Using early visual learning in a random dot motion coherence task, I expanded the knowledge about the link between spontaneous fluctuations in resting state activity and experience-dependant plasticity and learning. Participants in this project showed fast visual learning within one scanning session; the used task primarily activated visual and parietal brain areas. However, learning related changes in neural activity were present in the hippocampus. Crucially, even this rapid learning affected resting state dynamics both immediately after the learning as well as 24 h later. More precisely, the hippocampus

changed its coupling with the striatum in a way that was best explained as a consolidation of early learning related changes.

In order to examine the causal influence of spontaneous fluctuations on perception, I proceeded with a case study of an individual with a unilateral right parietal damage and visual extinction. Patients with visual extinction following right-hemisphere damage sometimes see and sometimes miss stimuli in the left visual field, particularly when stimuli are presented simultaneously to both visual fields. Awareness of left visual field stimuli is associated with increased activity in bilateral parietal and frontal cortex. I used this knowledge to investigate the outstanding question of *why* patients see or miss these stimuli. In order to do so, I again used functional MRI. This time I was able to show that prestimulus activity affects perception in the context of visual extinction following stroke. Measuring prestimulus activity in stimulus-responsive cortical areas during an extinction paradigm allowed me to compare prestimulus activity on physically identical bilateral trials that either did or did not lead to visual extinction. The significantly increased activity prior to stimulus presentation was observed in two areas that were also activated by visual stimulation: the left calcarine sulcus and right occipital inferior cortex. Furthermore, I established that effective connectivity within and between these areas was enhanced prior to stimulus presentation for bilateral seen trials. In summary, I provided evidence for the idea that differences in ongoing neural activity in visually responsive areas prior to stimulus onset affect awareness in visual extinction, and that these differences are mediated by fluctuations in extrinsic and intrinsic connectivity.

Having established the causal link between prestimulus activity and subsequent perception in visual extinction, I looked at the predictive power of prestimulus activity from a different and temporally more fine-grained angle using EEG and spectral analysis. Employing a similar random dot motion coherence task than in the first study, I demonstrated that alpha, beta and gamma oscillations seem to indicate a very specific brain state of "readiness" associated with the detection of coherent motion of a participant. By sparsely arranging stimuli every few seconds while recording scalp EEG, I was able to decouple each stimulus from any systematic effects of the preceding one and allowed for spontaneous fluctuations to occur. My analysis revealed that spontaneous fluctuations in ongoing activity are indeed correlated with the reported perceptual outcome of ambiguous stimulus processing. The increased power of the different oscillations bands prior to stimulus presentation exhibit a specific temporal and spatial pattern predicting motion detection on a single trial basis.

Finally, I investigated the link between a specific behavioural phenotype related to creativity, creative task performance and mind-wandering on the behavioural level, and validated these constructs by showing that the trait aspect of creativity is associated with physiological coupling between the DMN and a task-active network stimulated during a standard incubation task. Mind-wandering had previously been indicated as an "activity" that stimulates activity in opposing intrinsic brain networks, namely the DMN and task-active regions. Therefore, studying this omnipresent behaviour of humans seems to allow bridging the two worlds of ongoing and task-related brain activity. Pairing the incubation task with the investigation of trait and state measurements of creativity I was able to gain insight

into the neurophysiological basis of the relation between creativity and mind-wandering – a link that had only been investigated on the behavioural level before.

In summary, I have advanced the understanding about the functional role of spontaneous activity in general and its role in contextual changes in particular. Not only can learning – occurring as quickly as over one session – change resting-state activity and connectivity, but also is the spontaneous prestimulus activity predictive of behavioural outcomes in the lesioned and the healthy brain. Using EEG and fMRI, as well as spectral analyses and DCM to analyse connectivity, I showed that this predictive value manifests itself in a complex picture of oscillations and connectivity patterns. Finally, inter-individual differences in creativity and mind-wandering are also reflected in neuronal connectivity between networks implicated with an "idle" and an "active" brain.

7.2 Implications of this research

The main method used in this thesis is fMRI together with one EEG project. Therefore – within the scientific field studying ongoing brain activity – the implications drawn from the work presented here are focused on these two methods, with an emphasis on fMRI. In particular, it has contributed to the understanding of two aspects of ongoing activity, namely the observation that it predicts perception and that it is modulated by recent (learning) experiences as well as trait variables.

7.2.1 Ongoing activity predicts perception

Both, fMRI and EEG studies, have provided evidence for the idea that spontaneous activity predicts perception (on overview has been given in 1.3). This influence is

usually studied looking at the EEG or BOLD signal that precedes the actual presentation of the target stimulus, namely the prestimulus interval. Depending on the method used, this interval needs to be chosen with care and might be subject to inter- and intra-individual differences. There are no established standards yet as to how to define these intervals for any given technique or modality. Therefore, most studies seem to apply whatever works best in terms of their data. In Chapter 4 and Chapter 5 of this thesis, I present further data on two very different timescales that underline the importance of ongoing activity for subsequent behavioural perceptual outcomes. The results in Chapter 4 extend the findings observed in healthy brains and provide the first causal insight in the phenomenon of visual extinction. Further, the findings described in Chapter 5 provide the first detailed spectral description of the prestimulus interval preceding a random dot motion coherence task.

Previously, it has been argued about whether the interaction between ongoing and evoked brain activity is linear or not (Biswal et al., 1995; Senkowski et al., 2007; Becker et al., 2011; Schölvinck et al., 2012; He, 2013). Given the omnipresent nature of ongoing activity and the frequent ignorance of the majority of neuroscientists about it, I would like to emphasise the importance of taking it into account. Even though most neuroscientists have heard about spontaneous fluctuations of brain activity by now, there are no standard methods yet outlining how to control for it when studying task-evoked responses. Therefore, it is most often simply ignored. However, this might lead to wrong conclusions and is highly advised against.

7.2.2 Ongoing activity is modulated by learning and trait variables

The other direction of the two-way relationship between ongoing activity and behaviour, namely the potential of recent experiences such as learning as well as inter-individual trait differences to modulate ongoing activity and coupling patterns, was examined in Chapter 3 and Chapter 6. Chapter 3 provides the first demonstration that a (visual) learning task performed for as short as one session does not solely affect the activity in brain regions that are sensitive to the learning experience, but also leads to changes in the coupling between these regions. This change is preserved after 24 h, possibly excluding the possibility that the effect is merely a replay effect. This finding confirms the multifaceted pattern of variables that are all able to influence the pattern of ongoing activity present in our brain at any given moment.

Exploiting the observation that people who are more creative seem to mind-wander more, I established a three-way interaction between two behavioural and one neurophysiological measurement(s) providing evidence for the idea that individuals who are more creative exhibit a different pattern of neuronal coupling between two complementary intrinsic brain networks. These two findings contribute to the literature concerning the "formability" of ongoing activity due to short- and long-term behavioural measures, including learning (Albert, Robertson, and Miall, 2009; Smith et al., 2011) and traits like illness, mental diseases and addiction (Baliki et al., 2011; Liao et al., 2012; Yang and Tsai, 2013).

7.2.3 Cause and effect: the interplay between ongoing and evoked activity

Having established a bidirectional interaction between ongoing brain activity and behaviour in the widest sense, i.e. including phenotypic variations, one might start to wonder about the "chicken-and-egg question" about this relationship. In other words: if ongoing activity can be shaped by certain experiences, to what extent is it possible to exploit this effect such that one "trains" one's own activity to an extent that would in turn trigger a certain way of stimulus processing. This thought experiment creates an endless loop of interactions difficult – or impossible – to disentangle. In other words, the thought experiment asks which factors shape ongoing and evoked activity in the first place. Ultimately, the answer to this question might be about the nature versus nurture debate – where it has to be established to what extend our genes shape the patterns of ongoing activity in a (relatively[13]) predetermined way. This would need to be compared to the extent to which ongoing activity can be modified by experience or environmental influences respectively. Like for any other trait or behaviour that is supposedly influenced by "nature and nurture" a rather complicated entanglement might be observed, leaving the scientist once more with an unsolvable puzzle of interactions. Thus, one potential conclusion would be to simply give up. However, the other one is to take steps of reasonable size looking at parts of the puzzle by taking one piece at a time in order to gain further insight into the complex dynamics of ongoing brain activity and its manifestation on the behavioural level – keeping in mind that each single step is only part of the bigger puzzle.

[13] The word "relatively" is added here just for completeness, because epigenetic changes can be caused by environmental causes leading to modified gene expression.

Indeed, a small number of researchers have started to follow this direction and have for instance shown that ongoing activity also varies with age, and that this variation might correlate with changes in performance (Wenger et al., 2004; Colonnese and Khazipov, 2012; Vaden et al., 2012). Very recently, Harmelech and colleagues (2013) proposed a Hebbian-like rule for the changes in connectivity during rest after a single neurofeedback session. Similar to the changes I described in Chapter 3, they observed persistent changes after one day. Furthermore, these changes reflect the level of the regions' prior co-activation, i.e. during the neurofeedback session. Crucially, the effect was present on a single participant level. Such an observation might open the possibility to use ongoing activity as a kind of diagnostic window into an individual's history of prior brain activations.

7.3 Outstanding questions and conclusion

In this thesis, I have touched upon outstanding issues in the field of ongoing brain activity. The presented results contribute to the field, but also stimulate further experiments to explore these themes in the near future.

7.3.1 Timescale of changes in ongoing activity and its relation to structural changes

In Chapter 3 I showed that changes in the coupling between two brain regions in response to a visual learning experience as short as one session were present directly after the session during rest and were pertained 24 h after the experiment. To the best of my knowledge, this was the first study to report changes in coupling in ongoing activity at two different time points post-experiment. Only a very few studies have

started to explore the question of changes in ongoing activity in response to (learning) experience (see 1.3 for a summary of these), and even fewer have looked at long-term changes, i.e. the temporal profile of changes in resting state activity (Harmelech et al., 2013). Probably the most studied and therefore advanced field with regard to long-term changes in the human brain is the field of learning, more specifically motor learning. Indeed, it is known that functional, as well as structural (Scholz et al., 2009; Johansen-Berg, 2010; Tomassini et al., 2011) brain networks change after a new skill has been acquired. However, it is largely unknown if and how the rapid functional changes observed in the task related activity might sustain longer term changes in structure. In other words: what is the relationship between short-term and long-term motor memory in the human brain? Although the aforementioned role of resting state functional networks has been demonstrated previously (Miall and Robertson, 2006; Albert, Robertson, and Miall, 2009), a potential link between functional and structural changes has only been proposed recently (Taubert et al., 2011; Vahdat et al., 2011). In particular, Taubert et al. (2011) provided evidence that the learning of a complex motor task leads to truly long-lasting changes in functional resting-state networks. Even after six weeks, changes and further modulations of the fronto-parietal network connectivity in accordance with individual performance improvements were reported. These changes seem to be tightly correlated with the structural changes observed in grey matter. Thus, ongoing brain activity seems to be functionally relevant for morphological adaptions in the human brain.

Interestingly, this relationship appears to be – just like the one between ongoing activity and behaviour (as described in 7.2.3) – bidirectional: structural properties,

including the length, number, and spatial location of white matter streamlines are indicative of and can be inferred from the strength of resting-state (as well as task-based) functional correlations between different brain regions (Hermundstad et al., 2013).

Taken together, the temporal nature of changes in ongoing activity and its influence – or predictive role – for subsequent structural changes have just begun to be investigated. However, the increasing number of more recent studies taking both adaptive mechanisms – namely in functional and structural data – into account is promising to provide further insights in the near future.

7.3.2 Origin and scale of ongoing activity

In this thesis, I have investigated ongoing activity from two complementary perspectives using EEG and fMRI. This provided insights into its functional role on two different time-scales and allowed me to make inferences about relevant spectral, temporal and spatial patterns. Using fMRI and BOLD in three of the four studies described here, I focused on ultra-slow fluctuations in spontaneous activity in areas of the cerebral cortex. These fluctuations are dominated by very low temporal frequencies that follow a 1/f-like power distribution (Cordes et al., 2001). These fluctuations are also present in electrophysiological recordings, especially in gamma-band local field potential power (Schölvinck et al., 2010). The limited temporal resolution of fMRI can be addressed by using EEG or electrophysiological recordings to examine a potential overlap between spectral, spatial and temporal patterns. Here, I demonstrated a rather complex spectral pattern that can be involved in prestimulus activity predicting the behavioural outcome in a random dot motion

coherence task. Using an almost identical task in an fMRI paradigm revealed an equally complex picture of adaptations in coupling between brain regions involved in the learning process. But what is the relation between spontaneous brain activity at these different temporal and spatial scales?

Electrophysiological data of single neurons demonstrate that network activity on much smaller temporal and spatial scales can be shaped by spontaneous activity as well. For instance, the spontaneous firing of single neurons is correlated to the instantaneous spatial pattern of spontaneous population activity in neighboured neurons (Tsodyks et al., 1999). These spontaneously emerging spatial patterns in population activity in visual cortex resemble the well-known orientation maps (Kenet et al., 2003) suggesting that they might be shaped by intra-cortical connectivity. Taking this finding one step further, it might be the case that neurons in visual cortex fire to a certain degree in accordance with spontaneously emerging states – instead of simply responding to visual input from outside. In conclusion, the ongoing activity of single neurons and neuronal populations measured with EEG, as well as the haemodynamic activity measured with the BOLD contrast using fMRI, have a stable spatio-temporal structure – a finding that should prompt future studies in the field.

Having started to establish the relationship between structural and functional changes on a rather large scale as outlined in the previous paragraph, the connection between fast switching spontaneous cortical states described here and the ultra-slow fluctuations observed with fMRI remains obscure. Further research is needed to determine whether they share a common origin. Answering this question will include studies that trigger changes in the spontaneous activity on both scales, probably

involving behavioural measures related to the functional relevance of ongoing activity.

7.3.3 Conclusion or: The function of ongoing activity

The core interest of this thesis is the functional role of ongoing activity and I have already discussed its dynamics in detail, especially in light of the research questions I addressed in Chapter 3. I want to finish with a more general summary-like description of what research into ongoing activity has revealed so far about its function – emphasising that the brain seems to be doing "a lot" when it was traditionally assumed that it was doing "nothing at all" when observed at "rest". In other words, the observed activity during rest – when the mind usually starts to wander off – is not simply due to control mechanisms of automatic functions such as breathing and heart rate. In fact, the active networks appear to be quite similar to those observed during task performance. Thus, by studying these networks – that seem to consume the major share of energy used by the brain (Raichle and Mintun, 2006) – one can discover brain connectivity and how this might change in response to certain influences, such as age or disease. The mere observation, that ongoing activity is constantly present *but* modifiable suggests that it has to be important somehow.

However, what is the ultimate function of this ubiquitous activity? Several potential answers have been proposed. First, it might be one of the processes involved in memory consolidation (Albert, Robertson, Mehta, et al., 2009). Given its dynamic nature in response to certain learning experiences, it might be part of the neuronal correlate of so-called wake-dependant changes. These are involved in the

consolidation process, which includes sleep-dependent changes as well. Several animal studies have provided supportive evidence for the idea of an immediate replay of sensory experience in visual cortex and hippocampus (Foster and Wilson, 2006; Han et al., 2008). These patterns of recent sensory experiences in subsequent spontaneous activity might help to consolidate experiences into stable cortical modifications by contributing to short-term memory processes. Thus, understanding *how* ongoing activity is affected by learning might provide insight into the mechanisms and pathways responsible for determining in what way memories are consolidated during wakefulness, as well as during sleep.

Second, ongoing activity could serve some kind of maintenance function that helps to keep the brain's connections running when these are not in direct use, comparable to a car whose engine should be started once in a while to avoid failure. Similarly, it might help to maintain relationships between areas that are often active together to perform certain tasks, e.g. visual, auditory or cognitive tasks. In other words, the ongoing activity might help to organise and coordinate neural activity via supporting cortical network structures (Salinas and Sejnowski, 2001; Buzsáki and Draguhn, 2004).

A third suggestion was related – I use the past tense deliberately, because it is no longer considered a valid suggestion – to the quest for the neuronal correlates of consciousness. In the very early days of the study of ongoing activity, Greicius et al. (2003) suspected that they would find insight into the neural underpinnings of a critical but poorly understood component of human consciousness variably referred to as "a conscious resting state" during which real-time ongoing conscious

processing occurs. However, the finding that the same networks of ongoing activity are also present in altered states of consciousness such as sleep and anaesthesia led to the same researchers discarding this idea about a link between ongoing activity and conscious processing (Greicius et al., 2008; Boly et al., 2009; Vanhaudenhuyse et al., 2010).

Yet another alternative proposed for the function of ongoing activity is related to the theory of predictive coding (see 1.2.1.2). In this framework, ongoing activity helps to prime the brain to respond to future (sensory) stimuli (Engel et al., 2001; Pouget et al., 2003). Instead of an idle brain that is waiting for some input, the supporters of the predictive coding framework propose that the cycling activity in the networks exhibiting spontaneous fluctuations might help the brain to utilise past experiences to inform upcoming decisions, such as the perception of coherently moving dots. Electrophysiological recordings in cat visual cortex show that ongoing activity can predict subsequent neural responses to stimuli and the synchronous fluctuations between pairs of neurons observed prior to stimulus onset are indicative for the subsequent response latency (Fries et al., 2001). In favour of the predictive coding hypothesis, supporters have proposed that such a mode of operation would save a lot of energy. It is computationally demanding to calculate everything on the fly. Thus, the use of ongoing patterns that are guessing what might happen next can reduce the computational demand posed on the system. Furthermore, it can change perception as shown in the work presented here, for example. In other words, the brain might try to minimise surprise in order to be prepared for what occurs or happens next.

Of course, the possibility that all ongoing activity is simply an emergent property – a by-product – of the brain being alive remains. In other words, it might simply be there because of electric currents due to the brain being alive. These currents are naturally restricted by the anatomical connections providing them with a non-random structure.

In any case, it might be the case that all of these hypotheses are valid or true to a certain extent – because they are not mutually exclusive. The work presented here provided support for the first and fourth emphasising the bidirectional connection between ongoing activity and behaviour – where both can modify each other. I want to finish with a quote from Chris Miall, a researcher involved in the study of the dynamics of ongoing brain activity: "Whatever resting activity is doing, its existence certainly proves one thing: The brain only rests when you're dead." (Smith, 2012)

References

Acs, F., Greenlee, M.W., 2008. Connectivity modulation of early visual processing areas during covert and overt tracking tasks. NeuroImage 41, 380–8.

Adab, H.Z., Vogels, R., 2011. Practicing coarse orientation discrimination improves orientation signals in macaque cortical area v4. Current Biology 21, 1661–6.

Albert, N.B., Robertson, E.M., Mehta, P., Miall, R.C., 2009. Resting state networks and memory consolidation. Communicative & integrative biology 2, 530–2.

Albert, N.B., Robertson, E.M., Miall, R.C., 2009. The resting human brain and motor learning. Current biology : CB 19, 1023–7.

Allik, J., Rauk, M., Luuk, A., 1981. Control and sense of eye movement behind closed eyelids. Perception 10, 39–51.

Anticevic, A., Cole, M.W., Murray, J.D., Corlett, P.R., Wang, X.-J., Krystal, J.H., 2012. The role of default network deactivation in cognition and disease. Trends in cognitive sciences 16, 584–92.

Antrobus, J.S., 1968. Information theory and stimulus-independant thought. Br. J. Psychol. 59, 423–30.

Arieli, a, Sterkin, a, Grinvald, a, Aertsen, a, 1996. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. Science (New York, N.Y.) 273, 1868–71.

Arnal, L.H., Giraud, A.-L., 2012. Cortical oscillations and sensory predictions. Trends in cognitive sciences 16, 390–8.

Arnal, L.H., Wyart, V., Giraud, A.-L., 2011. Transitions in neural oscillations reflect prediction errors generated in audiovisual speech. Nature neuroscience 14, 797–801.

Attwell, D., Iadecola, C., 2002. The neural basis of functional brain imaging signals. Trends in neurosciences 25, 621–5.

Attwell, D., Laughlin, S.B., 2001. An energy budget for signaling in the grey matter of the brain. Journal of cerebral blood flow and metabolism : official journal of the International Society of Cerebral Blood Flow and Metabolism 21, 1133–45.

Azouz, R., Gray, C.M., 2008. Stimulus-selective spiking is driven by the relative timing of synchronous excitation and disinhibition in cat striate neurons in vivo. The European journal of neuroscience 28, 1286–300.

Baddeley, A.D., Hitch, G., 1993. The recency effect: Implicit learning with explicit retrieval? Memory & Cognition 21, 146–155.

Baird, B., Smallwood, J., Mrazek, M.D., Kam, J.W.Y., Franklin, M.S., Schooler, J.W., 2012. Inspired by distraction: mind wandering facilitates creative incubation. Psychological science 23, 1117–22.

Baker, C.L., Hess, R.F., Zihl, J., 1991. Residual motion perception in a "motion-blind" patient, assessed with limited-lifetime random dot stimuli. The Journal of Neuroscience 11, 454–61.

Baldassarre, A., Lewis, C.M., Committeri, G., Snyder, A.Z., 2012. Individual variability in functional connectivity predicts performance of a perceptual task. Proceedings of the National Academy of Sciences of the United States of America 109, 3516–21.

Baliki, M.N., Baria, A.T., Apkarian, A.V., 2011. The cortical rhythms of chronic back pain. The Journal of neuroscience : the official journal of the Society for Neuroscience 31, 13981–90.

Başar, E., Schürmann, M., Demiralp, T., Başar-Eroglu, C., Ademoglu, A., 2001. Event-related oscillations are "real brain responses"--wavelet analysis and new strategies. International journal of psychophysiology : official journal of the International Organization of Psychophysiology 39, 91–127.

Battaglia, F.P., Benchenane, K., Sirota, A., Pennartz, C.M.A., Wiener, S.I., 2011. The hippocampus: hub of brain network communication for memory. Trends in cognitive sciences 15, 310–8.

Baylis, G.C., Driver, J., Rafal, R.D., 1984. Visual Extinction and Stimulus Repetition. Journal of Cognitive Neuroscience 5, 453–66.

Beck, A.T., Steer, R.A., Brown, G.K., 1996. Manual for the Beck Depression Inventory-II. Psychological Corporation, San Antonio, Texas.

Becker, R., Reinacher, M., Freyer, F., Villringer, A., Ritter, P., 2011. How ongoing neuronal oscillations account for evoked fMRI variability. The Journal of neuroscience : the official journal of the Society for Neuroscience 31, 11016–27.

Beckmann, C.F., DeLuca, M., Devlin, J.T., Smith, S.M., 2005. Investigations into resting-state connectivity using independent component analysis. Philosophical transactions of the Royal Society of London. Series B, Biological sciences 360, 1001–13.

Berg, P., Scherg, M., 1994. A multiple source approach to the correction of eye artifacts. Electroencephalography and clinical neurophysiology 90, 229–41.

Berger, H., 1929. Über das Elektroenkephalogram des Menschen. Arch f Psychiat 87, 527–70.

Berti, A., Rizzolatti, G., Umana, F., 1987. Visual Processing without Awareness : Evidence from Unilateral Neglect. Journal of Cognitive Neuroscience 4, 345–51.

Bianciardi, M., Fukunaga, M., van Gelderen, P., Horovitz, S.G., de Zwart, J.A., Duyn, J.H., 2009. Modulation of spontaneous fMRI activity in human visual cortex by behavioral state. NeuroImage 45, 160–8.

Bianciardi, M., Fukunaga, M., van Gelderen, P., Horovitz, S.G., de Zwart, J.A., Shmueli, K., Duyn, J.H., 2009. Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: a 7 T study. Magnetic resonance imaging 27, 1019–29.

Birn, R.M., 2007. The behavioral significance of spontaneous fluctuations in brain activity. Neuron 56, 8–9.

Birn, R.M., Diamond, J.B., Smith, M.A., Bandettini, P.A., 2006. Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. NeuroImage 31, 1536–48.

Birn, R.M., Smith, M.A., Jones, T.B., Bandettini, P.A., 2008. The respiration response function: the temporal dynamics of fMRI signal fluctuations related to changes in respiration. NeuroImage 40, 644–54.

Bishop, G., 1933. Cyclic changes in excitability of the optic path- way of the rabbit. Am J Physio 103, 213–224.

Bishop, S.J., 2009. Trait anxiety and impoverished prefrontal control of attention. Nature neuroscience 12, 92–8.

Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine 34, 537–41.

Blakemore, B.C., Campbell, F.W., 1969. On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. Journal of Physiology 203, 237–260.

Bollimunta, A., Chen, Y., Schroeder, C.E., Ding, M., 2008. Neuronal mechanisms of cortical alpha oscillations in awake-behaving macaques. The Journal of neuroscience : the official journal of the Society for Neuroscience 28, 9976–88.

Boly, M., Balteau, E., Schnakers, C., Degueldre, C., Moonen, G., Luxen, a, Phillips, C., Peigneux, P., Maquet, P., Laureys, S., 2007. Baseline brain activity

fluctuations predict somatosensory perception in humans. Proceedings of the National Academy of Sciences of the United States of America 104, 12187–92.

Boly, M., Tshibanda, L., Vanhaudenhuyse, A., Noirhomme, Q., Schnakers, C., Ledoux, D., Boveroux, P., Garweg, C., Lambermont, B., Phillips, C., Luxen, A., Moonen, G., Bassetti, C., Maquet, P., Laureys, S., 2009. Functional connectivity in the default network during resting state is preserved in a vegetative but not in a brain dead patient. Human brain mapping 30, 2393–400.

Breakspear, M., 2004. "Dynamic" connectivity in neural systems: theoretical and empirical considerations. Neuroinformatics 2, 205–26.

Brooks, J.L., 2012. Counterbalancing for serial order carryover effects in experimental condition orders. Psychological methods 17, 600–14.

Brovelli, A., Laksiri, N., Nazarian, B., Meunier, M., Boussaoud, D., 2008. Understanding the neural computations of arbitrary visuomotor learning through fMRI and associative learning theory. Cerebral cortex (New York, N.Y. : 1991) 18, 1485–95.

Buckley, M.J., 2005. The role of the perirhinal cortex and hippocampus in learning, memory, and perception. Q J Exp Psychol B 58, 246–68.

Buckner, R.L., Andrews-Hanna, J.R., Schacter, D.L., 2008. The brain's default network: anatomy, function, and relevance to disease. Annals of the New York Academy of Sciences 1124, 1–38.

Busch, N.A., Dubois, J., VanRullen, R., 2009. The phase of ongoing EEG oscillations predicts visual perception. The Journal of neuroscience : the official journal of the Society for Neuroscience 29, 7869–76.

Buxton, R.B., Wong, E.C., Frank, L.R., 1998. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine 39, 855–64.

Buzsáki, G., Draguhn, A., 2004. Neuronal oscillations in cortical networks. Science (New York, N.Y.) 304, 1926–9.

Carson, S.H., Peterson, J.B., Higgins, D.M., 2005. Reliability, Validity, and Factor Structure of the Creative Achievement Questionnaire. Creativity Research Journal 17, 37–50.

Chavan, C.F., Manuel, A.L., Mouthon, M., Spierer, L., 2013. Spontaneous pre-stimulus fluctuations in the activity of right fronto-parietal areas influence inhibitory control performance. Frontiers in human neuroscience 7, 238.

Chen, C.C., Henson, R.N., Stephan, K.E., Kilner, J.M., Friston, K.J., 2009. Forward and backward connections in the brain: a DCM study of functional asymmetries. NeuroImage 45, 453–62.

Chen, J.J., Pike, G.B., 2009. Origins of the BOLD post-stimulus undershoot. NeuroImage 46, 559–68.

Chen, S., Ross, T.J., Zhan, W., Myers, C.S., Chuang, K.-S., Heishman, S.J., Stein, E.A., Yang, Y., 2008. Group independent component analysis reveals consistent resting-state networks across multiple sessions. Brain research 1239, 141–51.

Chiu, C., Weliky, M., 2001. Spontaneous activity in developing ferret visual cortex in vivo. The Journal of neuroscience : the official journal of the Society for Neuroscience 21, 8906–14.

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. Proceedings of the National Academy of Sciences of the United States of America 106, 8719–24.

Ciçek, M., Gitelman, D., Hurley, R.S.E., Nobre, A., Mesulam, M., 2007. Anatomical physiology of spatial extinction. Cerebral Cortex 17, 2892–8.

Cohen, J., 1988. Statistical Power Analysis for the Behavioral Sciences, 2nd ed. Academic Press, New York.

Cole, D.M., Smith, S.M., Beckmann, C.F., 2010. Advances and pitfalls in the analysis and interpretation of resting-state FMRI data. Frontiers in systems neuroscience 4, 8.

Colgin, L.L., Moser, E.I., Moser, M.-B., 2008. Understanding memory through hippocampal remapping. Trends in Neurosciences 31, 469–77.

Collewijn, H., van der Steen, J., Steinman, R.M., 1985. Human eye movements associated with blinks and prolonged eyelid closure. Journal of neurophysiology 54, 11–27.

Colonnese, M., Khazipov, R., 2012. Spontaneous activity in developing sensory circuits: Implications for resting state fMRI. NeuroImage 62, 2212–21.

Corbetta, M., 2012. Functional connectivity and neurological recovery. Developmental psychobiology 54, 239–53.

Corbetta, M., Kincade, J.M., Shulman, G.L., 2002. Neural systems for visual orienting and their relationships to spatial working memory. Journal of cognitive neuroscience 14, 508–23.

Cordes, D., Haughton, V.M., Arfanakis, K., Carew, J.D., Turski, P.A., Moritz, C.H., Quigley, M.A., Meyerand, M.E., 2001. Frequencies Contributing to Functional Connectivity in the Cerebral Cortex in "Resting-state" Data. AJNR Am. J. Neuroradiol. 22, 1326–1333.

Cordes, D., Haughton, V.M., Arfanakis, K., Wendt, G.J., Turski, P.A., Moritz, C.H., Quigley, M.A., Meyerand, M.E., 2000. Mapping functionally related regions of brain with functional connectivity MR imaging. AJNR. American journal of neuroradiology 21, 1636–44.

Costa-Faidella, J., Baldeweg, T., Grimm, S., Escera, C., 2011. Interactions between "what" and "when" in the auditory system: temporal predictability enhances repetition suppression. The Journal of neuroscience : the official journal of the Society for Neuroscience 31, 18590–7.

Coste, C.P., Sadaghiani, S., Friston, K.J., Kleinschmidt, A., 2011. Ongoing brain activity fluctuations directly account for intertrial and indirectly for intersubject variability in Stroop task performance. Cerebral Cortex 21, 2612–9.

Cowey, A., Marcar, V.L., 1992. The Effect of Removing Superior Temporal Cortical Motion Areas in the Macaque Monkey: I. Motion Discrimination Using Simple Dots. The European Journal of Neuroscience 4, 1219–1227.

Damoiseaux, J.S., Rombouts, S.A.R.B., Barkhof, F., Scheltens, P., Stam, C.J., Smith, S.M., Beckmann, C.F., 2006. Consistent resting-state networks across healthy subjects. Proceedings of the National Academy of Sciences of the United States of America 103, 13848–53.

David, O., Guillemain, I., Saillet, S., Reyt, S., Deransart, C., Segebarth, C., Depaulis, A., 2008. Identifying neural drivers with functional MRI: an electrophysiological validation. PLoS biology 6, 2683–97.

De Luca, M., Smith, S., De Stefano, N., Federico, A., Matthews, P.M., 2005. Blood oxygenation level dependent contrast resting state networks are relevant to functional activity in the neocortical sensorimotor system. Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale 167, 587–94.

Deco, G., Jirsa, V.K., McIntosh, A.R., 2011. Emerging concepts for the dynamical organization of resting-state activity in the brain. Nature reviews. Neuroscience 12, 43–56.

Deco, G., Romo, R., 2008. The role of fluctuations in perception. Trends in neurosciences 31, 591–8.

Dijksterhuis, A., Meurs, T., 2006. Where creativity resides: the generative power of unconscious thought. Consciousness and cognition 15, 135–46.

Ding, L., Gold, J.I., 2010. Caudate encodes multiple computations for perceptual decisions. The Journal of Neuroscience 30, 15747–59.

Donner, T.H., Siegel, M., 2011. A framework for local cortical oscillation patterns. Trends in cognitive sciences 15, 191–9.

Donner, T.H., Siegel, M., Fries, P., Engel, A.K., 2009. Buildup of choice-predictive activity in human motor cortex during perceptual decision making. Current biology : CB 19, 1581–5.

Donner, T.H., Siegel, M., Oostenveld, R., Fries, P., Bauer, M., Engel, A.K., 2007. Population activity in the human dorsal pathway predicts the accuracy of visual motion detection. Journal of neurophysiology 98, 345–59.

Doucet, G., Naveau, M., Petit, L., Zago, L., Crivello, F., Jobard, G., Delcroix, N., Mellet, E., Tzourio-Mazoyer, N., Mazoyer, B., Joliot, M., 2012. Patterns of hemodynamic low-frequency oscillations in the brain are modulated by the nature of free thought during rest. NeuroImage 59, 3194–200.

Driver, J., Vuilleumier, P., 2001. Perceptual awareness and its loss in unilateral neglect and extinction. Cognition 79, 39–88.

Driver, J., Vuilleumier, P., Eimer, M., Rees, G., 2001. Functional magnetic resonance imaging and evoked potential correlates of conscious and unconscious vision in parietal extinction patients. NeuroImage 14, S68–75.

Dudai, Y., 2012. The restless engram: consolidations never end. Annual Review of Neuroscience 35, 227–47.

Ebbinghaus, H., 1913. On memory: A contribution to experimental psychology. Teachers College, New York.

Eickhoff, S.B., Stephan, K.E., Mohlberg, H., Grefkes, C., Fink, G.R., Amunts, K., Zilles, K., 2005. A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data. NeuroImage 25, 1325–35.

Ellwood, S., Pallier, G., Snyder, A., Gallate, J., 2009. The incubation effect: Hatching a solution? Creativity Research Journal 21, 6–14.

Engel, A.K., Fries, P., 2010. Beta-band oscillations--signalling the status quo? Current opinion in neurobiology 20, 156–65.

Engel, A.K., Fries, P., Singer, W., 2001. Dynamic predictions: oscillations and synchrony in top-down processing. Nature reviews. Neuroscience 2, 704–16.

Ermentrout, G.B., Galán, R.F., Urban, N.N., 2008. Reliability, synchrony and noise. Trends in neurosciences 31, 428–34.

Esterman, M., Noonan, S.K., Rosenberg, M., Degutis, J., 2012. In the Zone or Zoning Out? Tracking Behavioral and Neural Fluctuations During Sustained Attention. Cerebral cortex (New York, N.Y. : 1991).

Fair, D.A., Cohen, A.L., Dosenbach, N.U.F., Church, J.A., Miezin, F.M., Barch, D.M., Raichle, M.E., Petersen, S.E., Schlaggar, B.L., 2008. The maturing architecture of the brain's default network. Proceedings of the National Academy of Sciences of the United States of America 105, 4028–32.

Feldman, H., Friston, K.J., 2010. Attention, uncertainty, and free-energy. Frontiers in Human Neuroscience 4, 215.

Fiorentini, A., Berardi, N., 1981. Learning in grating waveform discrimination: specificity for orientation and spatial frequency. Vision Research 21, 1149–58.

Fiser, J., Chiu, C., Weliky, M., 2004. Small modulation of ongoing cortical dynamics by sensory input during natural vision. Nature 431, 573–8.

Forkstam, C., Petersson, K.M., 2005. Towards an explicit account of implicit learning. Current opinion in neurology 18, 435–41.

Forstmann, B.U., Anwander, A., Schäfer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., Bogacz, R., Turner, R., 2010. Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. Proceedings of the National Academy of Sciences of the United States of America 107, 15916–20.

Fortin, N.J., Agster, K.L., Eichenbaum, H.B., 2002. Critical role of the hippocampus in memory for sequences of events. Nature Neuroscience 5, 458–62.

Foster, D.J., Wilson, M.A., 2006. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature 440, 680–3.

Fox, M.D., Raichle, M.E., 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nature reviews. Neuroscience 8, 700–11.

Fox, M.D., Snyder, A.Z., Vincent, J.L., Corbetta, M., Van Essen, D.C., Raichle, M.E., 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. Proceedings of the National Academy of Sciences of the United States of America 102, 9673–8.

Fox, M.D., Snyder, A.Z., Vincent, J.L., Raichle, M.E., 2007. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. Neuron 56, 171–84.

Fox, M.D., Snyder, A.Z., Zacks, J.M., Raichle, M.E., 2006. Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. Nature neuroscience 9, 23–5.

Fox, M.D., Zhang, D., Snyder, A.Z., Raichle, M.E., 2009. The global signal and observed anticorrelated resting state brain networks. Journal of neurophysiology 101, 3270–83.

Fox, P.T., Raichle, M.E., Mintun, M.A., Dence, C., 1988. Nonoxidative glucose consumption during focal physiologic neural activity. Science (New York, N.Y.) 241, 462–4.

Franco, A.R., Pritchard, A., Calhoun, V.D., Mayer, A.R., 2009. Interrater and intermethod reliability of default mode network selection. Human brain mapping 30, 2293–303.

Fransson, P., 2005. Spontaneous low-frequency BOLD signal fluctuations: an fMRI investigation of the resting-state default mode of brain function hypothesis. Human brain mapping 26, 15–29.

Frensch, P.A., Runger, D., 2003. Implicit learning. Current Directions in Psychological Science 12, 13–18.

Freud, S., 1909a. Analysis of a phobia of a five year old boy, in: The Pelican Freud Library (1977). pp. 169–306.

Freud, S., 1909b. Bemerkungen über einen Fall von Zwangsneurose (Der "Rattenmann"). Jb. psychoanal. psychopathol. Forsch. 357-421.

Fries, P., 2009. Neuronal gamma-band synchronization as a fundamental process in cortical computation. Annual review of neuroscience 32, 209–24.

Fries, P., Neuenschwander, S., Engel, A.K., Goebel, R., Singer, W., 2001. Rapid feature selective neuronal synchronization through correlated latency shifting. Nature neuroscience 4, 194–200.

Friston, K., 2005. A theory of cortical responses. Philosophical transactions of the Royal Society of London. Series B, Biological sciences 360, 815–36.

Friston, K., 2012. Ten ironic rules for non-statistical reviewers. NeuroImage 61, 1300–10.

Friston, K., Stephan, K.E., Li, B., Daunizeau, J., 2010. Generalised filtering. Math. Probl. Eng. 621670.

Friston, K.J., 1994. Functional and effective connectivity in neuroimaging: A synthesis. Human Brain Mapping 2, 56–78.

Friston, K.J., 2002. Bayesian estimation of dynamical systems: an application to fMRI. NeuroImage 16, 513–30.

Friston, K.J., 2011. Functional and effective connectivity: a review. Brain connectivity 1, 13–36.

Friston, K.J., Buechel, C., Fink, G.R., Morris, J., Rolls, E., Dolan, R.J., 1997. Psychophysiological and modulatory interactions in neuroimaging. NeuroImage 6, 218–29.

Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. NeuroImage 19, 1273–302.

Friston, K.J., Mechelli, A., Turner, R., Price, C.J., 2000. Nonlinear responses in fMRI: the Balloon model, Volterra kernels, and other hemodynamics. NeuroImage 12, 466–77.

Fujioka, T., Trainor, L.J., Large, E.W., Ross, B., 2009. Beta and gamma rhythms in human auditory cortex during musical beat processing. Annals of the New York Academy of Sciences 1169, 89–92.

Fukunaga, M., Horovitz, S.G., van Gelderen, P., de Zwart, J.A., Jansma, J.M., Ikonomidou, V.N., Chu, R., Deckers, R.H.R., Leopold, D.A., Duyn, J.H., 2006. Large-amplitude, spatially correlated fluctuations in BOLD fMRI signals during extended rest and early sleep stages. Magnetic resonance imaging 24, 979–92.

Galéra, C., Orriols, L., M'Bailara, K., Laborey, M., Contrand, B., Ribéreau-Gayon, R., Masson, F., Bakiri, S., Gabaude, C., Fort, A., Maury, B., Lemercier, C., Cours, M., Bouvard, M.-P., Lagarde, E., 2012. Mind wandering and driving: responsibility case-control study. BMJ (Clinical research ed.) 345, e8105.

Garrido, M.I., Kilner, J.M., Kiebel, S.J., Stephan, K.E., Friston, K.J., 2007. Dynamic causal modelling of evoked potentials: a reproducibility study. NeuroImage 36, 571–80.

Gavrilescu, M., Shaw, M.E., Stuart, G.W., Eckersley, P., Svalbe, I.D., Egan, G.F., 2002. Simulation of the effects of global normalization procedures in functional MRI. NeuroImage 17, 532–42.

George, M.S., Mercer, J.S., Walker, R., Manly, T., 2008. A demonstration of endogenous modulation of unilateral spatial neglect: the impact of apparent time-pressure on spatial bias. Journal of the International Neuropsychological Society 14, 33–41.

Gerretsen, P., Graff-Guerrero, A., Menon, M., Pollock, B.G., Kapur, S., Vasdev, N., Houle, S., Mamo, D., 2010. Is desire for social relationships mediated by the serotonergic system in the prefrontal cortex? An [(18)F]setoperone PET study. Social neuroscience 5, 375–83.

Gerstein, G.L., 1960. Analysis of Firing Pafferns in Single Neurons. Science (New York, N.Y.) 131, 1811–2.

Gheysen, F., Van Opstal, F., Roggeman, C., Van Waelvelde, H., Fias, W., 2011. The neural basis of implicit perceptual sequence learning. Frontiers in Human Neuroscience 5, 137.

Giesbrecht, T., Jongen, E.M.M., Smulders, F.T.Y., Merckelbach, H., 2006. Dissociation, resting EEG, and subjective sleep experiences in undergraduates. The Journal of Nervous and Mental Disease 194, 362–8.

Glover, G.H., Li, T.Q., Ress, D., 2000. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine 44, 162–7.

Goebel, R., Roebroeck, A., Kim, D.-S., Formisano, E., 2003. Investigating directed cortical interactions in time-resolved fMRI data using vector autoregressive modeling and Granger causality mapping. Magnetic resonance imaging 21, 1251–61.

Goldman, R.I., Stern, J.M., Engel, J., Cohen, M.S., 2002. Simultaneous EEG and fMRI of the alpha rhythm. Neuroreport 13, 2487–92.

Goldstone, R.L., 1998. Perceptual learning. Annual review of psychology 49, 585–612.

Graham, K.S., Scahill, V.L., Hornberger, M., Barense, M.D., Lee, A.C.H., Bussey, T.J., Saksida, L.M., 2006. Abnormal categorization and perceptual learning in patients with hippocampal damage. The Journal of Neuroscience 26, 7547–54.

Greicius, M.D., Kiviniemi, V., Tervonen, O., Vainionpää, V., Alahuhta, S., Reiss, A.L., Menon, V., 2008. Persistent default-mode network connectivity during light sedation. Human brain mapping 29, 839–47.

Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. Proceedings of the National Academy of Sciences of the United States of America 100, 253–8.

Greicius, M.D., Supekar, K., Menon, V., Dougherty, R.F., 2009. Resting-state functional connectivity reflects structural connectivity in the default mode network. Cerebral cortex (New York, N.Y. : 1991) 19, 72–8.

Grol, M.J., Majdandzić, J., Stephan, K.E., Verhagen, L., Dijkerman, H.C., Bekkering, H., Verstraten, F.A.J., Toni, I., 2007. Parieto-frontal connectivity during visually guided grasping. The Journal of neuroscience : the official journal of the Society for Neuroscience 27, 11877–87.

Groppe, D.M., Urbach, T.P., Kutas, M., 2011. Mass univariate analysis of event-related brain potentials/fields I: a critical tutorial review. Psychophysiology 48, 1711–25.

Guilford, J.P., 1967. The nature of human intelligence. McGraw-Hill, New York.

Guilford, J.P., Christensen, P.R., Merrifield, P.R., Wilson, R.C., 1960. Alternate Uses. Sheridan Supply Co.

Haegens, S., Händel, B.F., Jensen, O., 2011. Top-down controlled alpha band activity in somatosensory areas determines behavioral performance in a discrimination task. The Journal of Neuroscience 31, 5197–204.

Hampson, M., Peterson, B.S., Skudlarski, P., Gatenby, J.C., Gore, J.C., 2002. Detection of functional connectivity using temporal correlations in MR images. Human brain mapping 15, 247–62.

Han, F., Caporale, N., Dan, Y., 2008. Reverberation of recent visual experience in spontaneous cortical waves. Neuron 60, 321–7.

Hanslmayr, S., Aslan, A., Staudigl, T., Klimesch, W., Herrmann, C.S., Bäuml, K.-H., 2007. Prestimulus oscillations predict visual perception performance between and within subjects. NeuroImage 37, 1465–73.

Harmelech, T., Preminger, S., Wertman, E., Malach, R., 2013. The day-after effect: long term, Hebbian-like restructuring of resting-state fMRI patterns induced by a single epoch of cortical activation. The Journal of Neuroscience 33, 9488–97.

Harrison, L., Penny, W.D., Friston, K., 2003. Multivariate autoregressive modeling of fMRI time series. NeuroImage 19, 1477–91.

Hasson, U., Nusbaum, H.C., Small, S.L., 2009. Task-dependent organization of brain regions active during rest. Proceedings of the National Academy of Sciences of the United States of America 106, 10841–6.

Hayasaka, S., Nichols, T.E., 2004. Combining voxel intensity and cluster extent with permutation test framework. NeuroImage 23, 54–63.

He, B.J., 2013. Spontaneous and task-evoked brain activity negatively interact. The Journal of neuroscience : the official journal of the Society for Neuroscience 33, 4672–82.

He, B.J., Raichle, M.E., 2009. The fMRI signal, slow cortical potential and consciousness. Trends in cognitive sciences 13, 302–9.

He, B.J., Snyder, A.Z., Zempel, J.M., Smyth, M.D., Raichle, M.E., 2008. Electrophysiological correlates of the brain's intrinsic large-scale functional

architecture. Proceedings of the National Academy of Sciences of the United States of America 105, 16039–44.

Heilman, K.M., Watson, R.T., Valenstein, E., 1994. Neglect: clincial and anatomical aspects, in: Feinberg, T.E., Farah, M.J. (Eds.), Behavioural Neurology and Neuropsychology. McGraw-Hill, New York, pp. 309–17.

Henry, G.H., Bishop, P.O., Tupper, R.M., Dreher, B., 1973. Orientation specificity and response variability of cells in the striate cortex. Vision research 13, 1771–9.

Hermundstad, A.M., Bassett, D.S., Brown, K.S., Aminoff, E.M., Clewett, D., Freeman, S., Frithsen, A., Johnson, A., Tipper, C.M., Miller, M.B., Grafton, S.T., Carlson, J.M., 2013. Structural foundations of resting-state and task-based functional connectivity in the human brain. Proceedings of the National Academy of Sciences of the United States of America 110, 6169–74.

Herrmann, C.S., Munk, M.H.J., Engel, A.K., 2004. Cognitive functions of gamma-band activity: memory match and utilization. Trends in cognitive sciences 8, 347–55.

Hesselmann, G., Kell, C. a, Eger, E., Kleinschmidt, A., 2008. Spontaneous local variations in ongoing neural activity bias perceptual decisions. Proceedings of the National Academy of Sciences of the United States of America 105, 10984–9.

Hesselmann, G., Kell, C. a, Kleinschmidt, A., 2008. Ongoing activity fluctuations in hMT+ bias the perception of coherent visual motion. The Journal of Neuroscience 28, 14481–5.

Hillis, A.E., Chang, S., Heidler-Gary, J., Newhart, M., Kleinman, J.T., Davis, C., Barker, P.B., Aldrich, E., Ken, L., 2006. Neural correlates of modality-specific spatial extinction. Journal of Cognitive Neuroscience 18, 1889–98.

Himberg, J., Hyvärinen, A., Esposito, F., 2004. Validating the independent components of neuroimaging time series via clustering and visualization. NeuroImage 22, 1214–22.

Hipp, J.F., Engel, A.K., Siegel, M., 2011. Oscillatory synchronization in large-scale cortical networks predicts perception. Neuron 69, 387–96.

Holcman, D., Tsodyks, M., 2006. The emergence of Up and Down states in cortical networks. PLoS computational biology 2, e23.

Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., Hagmann, P., 2009. Predicting human resting-state functional connectivity from structural connectivity. Proceedings of the National Academy of Sciences of the United States of America 106, 2035–40.

Horovitz, S.G., Braun, A.R., Carr, W.S., Picchioni, D., Balkin, T.J., Fukunaga, M., Duyn, J.H., 2009. Decoupling of the brain's default mode network during deep sleep. Proceedings of the National Academy of Sciences of the United States of America 106, 11376–81.

Horovitz, S.G., Fukunaga, M., de Zwart, J.A., van Gelderen, P., Fulton, S.C., Balkin, T.J., Duyn, J.H., 2008. Low frequency BOLD fluctuations during resting wakefulness and light sleep: a simultaneous EEG-fMRI study. Human brain mapping 29, 671–82.

Hutton, C., Josephs, O., Stadler, J., Featherstone, E., Reid, A., Speck, O., Bernarding, J., Weiskopf, N., 2011. The impact of physiological noise correction on fMRI at 7 T. NeuroImage 57, 101–12.

Irle, E., Ruhleder, M., Lange, C., Seidler-Brandler, U., Salzer, S., Dechent, P., Weniger, G., Leibing, E., Leichsenring, F., 2010. Reduced amygdalar and hippocampal size in adults with generalized social phobia. Journal of psychiatry & neuroscience : JPN 35, 126–31.

Iversen, J.R., Repp, B.H., Patel, A.D., 2009. Top-down control of rhythm perception modulates early auditory responses. Annals of the New York Academy of Sciences 1169, 58–73.

Johansen-Berg, H., 2010. Behavioural relevance of variation in white matter microstructure. Current opinion in neurology 23, 351–8.

Jones, M.R., Moynihan, H., MacKenzie, N., Puente, J., 2002. Temporal aspects of stimulus-driven attending in dynamic arrays. Psychological science 13, 313–9.

Kahnt, T., Grueschow, M., Speck, O., Haynes, J.-D., 2011. Perceptual learning and decision-making in human medial frontal cortex. Neuron 70, 549–59.

Kanai, R., Bahrami, B., Rees, G., 2010. Human parietal cortex structure predicts individual differences in perceptual rivalry. Current biology : CB 20, 1626–30.

Kanai, R., Rees, G., 2011. The structural basis of inter-individual differences in human behaviour and cognition. Nature reviews. Neuroscience 12, 231–42.

Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., Arieli, A., 2003. Spontaneously emerging cortical representations of visual attributes. Nature 425, 954–6.

Kersten, D., Mamassian, P., Yuille, A., 2004. Object perception as Bayesian inference. Annual review of psychology 55, 271–304.

Khader, P., Schicke, T., Röder, B., Rösler, F., 2008. On the relationship between slow cortical potentials and BOLD signal changes in humans. International

journal of psychophysiology : official journal of the International Organization of Psychophysiology 67, 252–61.

Killingsworth, M.A., Gilbert, D.T., 2010. A Wandering Mind Is an Unhappy Mind. Science 330, 932–932.

Kiviniemi, V., Kantola, J.-H., Jauhiainen, J., Hyvärinen, A., Tervonen, O., 2003. Independent component analysis of nondeterministic fMRI signal sources. NeuroImage 19, 253–60.

Kiviniemi, V., Starck, T., Remes, J., Long, X., Nikkinen, J., Haapea, M., Veijola, J., Moilanen, I., Isohanni, M., Zang, Y.-F., Tervonen, O., 2009. Functional segmentation of the brain cortex using high model order group PICA. Human brain mapping 30, 3865–86.

Kleinschmidt, A., Sterzer, P., Rees, G., 2012. Variability of perceptual multistability: from brain state to individual trait. Philosophical transactions of the Royal Society of London. Series B, Biological sciences 367, 988–1000.

Klimesch, W., 1999. EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. Brain research. Brain research reviews 29, 169–95.

Klimesch, W., Sauseng, P., Hanslmayr, S., Gruber, W., Freunberger, R., 2007. Event-related phase reorganization may explain evoked neural dynamics. Neuroscience and biobehavioral reviews 31, 1003–16.

Kosslyn, S.M., Intriligator, J.M., 1992. Is cognitive neuropsychology plausible? The perils of sitting on a one-legged stool. Journal of cognitive neuroscience 4, 96–105.

Kuba, M., Kubová, Z., Kremlácek, J., Langrová, J., 2007. Motion-onset VEPs: characteristics, methods, and diagnostic use. Vision research 47, 189–202.

Kumar, S., Stephan, K.E., Warren, J.D., Friston, K.J., Griffiths, T.D., 2007. Hierarchical processing of auditory objects in humans. PLoS computational biology 3, e100.

L, W., D F, H., I B, H., J, L., 2012. A systematic review of resting-state functional-MRI studies in major depression. Journal of affective disorders 142, 6–12.

Ladavas, E., Paladini, R., Cubellli, R., 1993. Implicit Associative Priming in a Patient with Left Visual Neglect. Neuropsychologia 31, 1307–20.

Lakatos, P., Karmos, G., Mehta, A.D., Ulbert, I., Schroeder, C.E., 2008. Entrainment of neuronal oscillations as a mechanism of attentional selection. Science (New York, N.Y.) 320, 110–3.

Larson-Prior, L.J., Zempel, J.M., Nolan, T.S., Prior, F.W., Snyder, A.Z., Raichle, M.E., 2009. Cortical network functional connectivity in the descent to sleep. Proceedings of the National Academy of Sciences of the United States of America 106, 4489–94.

Laufs, H., Krakow, K., Sterzer, P., Eger, E., Beyerle, A., Salek-Haddadi, A., Kleinschmidt, A., 2003. Electroencephalographic signatures of attentional and cognitive default modes in spontaneous brain activity fluctuations at rest. Proceedings of the National Academy of Sciences of the United States of America 100, 11053–8.

Law, C., Gold, J.I., 2009. Reinforcement learning can account for associative and perceptual learning on a visual decision task. Nature Neuroscience 12, 655–663.

Law, C.-T., Gold, J.I., 2008. Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nature Neuroscience 11, 505–13.

Lee, A.C.H., Bussey, T.J., Murray, E. a, Saksida, L.M., Epstein, R. a, Kapur, N., Hodges, J.R., Graham, K.S., 2005. Perceptual deficits in amnesia: challenging the medial temporal lobe "mnemonic" view. Neuropsychologia 43, 1–11.

Lenth, R.V., 2001. Some practical guidelines for effective sample-size. Determin. Am. Stat. 55, 187–93.

Leopold, D.A., Murayama, Y., Logothetis, N.K., 2003. Very slow activity fluctuations in monkey visual cortex: implications for functional brain imaging. Cerebral cortex (New York, N.Y. : 1991) 13, 422–33.

Levinson, D.B., Smallwood, J., Davidson, R.J., 2012. The persistence of thought: evidence for a role of working memory in the maintenance of task-unrelated thinking. Psychological science 23, 375–80.

Lewis, C.M., Baldassarre, A., Committeri, G., Romani, G.L., Corbetta, M., 2009. Learning sculpts the spontaneous activity of the resting human brain. Proceedings of the National Academy of Sciences of the United States of America 106, 17558–63.

Li, B., Daunizeau, J., Stephan, K.E., Penny, W., Hu, D., Friston, K., 2011. Generalised filtering and stochastic DCM for fMRI. NeuroImage 58, 442–57.

Liao, Y., Tang, J., Fornito, A., Liu, T., Chen, X., Chen, H., Xiang, X., Wang, X., Hao, W., 2012. Alterations in regional homogeneity of resting-state brain activity in ketamine addicts. Neuroscience letters 522, 36–40.

Lieberman, M.D., Chang, G.Y., Chiao, J., Bookheimer, S.Y., Knowlton, B.J., 2004. An event-related fMRI study of artificial grammar learning in a balanced chunk strength design. Journal of Cognitive Neuroscience 16, 427–38.

Linkenkaer-Hansen, K., Nikulin, V. V, Palva, S., Ilmoniemi, R.J., Palva, J.M., 2004. Prestimulus oscillations enhance psychophysical performance in humans. The Journal of neuroscience : the official journal of the Society for Neuroscience 24, 10186–90.

Logothetis, N.K., 2008. What we can do and what we cannot do with fMRI. Nature 453, 869–78.

Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. Nature 412, 150–7.

Lohrenz, T., McCabe, K., Camerer, C.F., Montague, P.R., 2007. Neural signature of fictive learning signals in a sequential investment task. Proceedings of the National Academy of Sciences of the United States of America 104, 9493–8.

Long, X.-Y., Zuo, X.-N., Kiviniemi, V., Yang, Y., Zou, Q.-H., Zhu, C.-Z., Jiang, T.-Z., Yang, H., Gong, Q.-Y., Wang, L., Li, K.-C., Xie, S., Zang, Y.-F., 2008. Default mode network as revealed with multiple methods for resting-state functional MRI analysis. Journal of neuroscience methods 171, 349–55.

Lopes da Silva, F.H., Vos, J.E., Mooibroek, J., Van Rotterdam, A., 1980. Relative contributions of intracortical and thalamo-cortical processes in the generation of alpha rhythms, revealed by partial coherence analysis. Electroencephalography and clinical neurophysiology 50, 449–56.

Lowe, M.J., Mock, B.J., Sorenson, J.A., 1998. Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. NeuroImage 7, 119–32.

Lu, H., Zuo, Y., Gu, H., Waltz, J.A., Zhan, W., Scholl, C.A., Rea, W., Yang, Y., Stein, E.A., 2007. Synchronized delta oscillations correlate with the resting-state functional MRI signal. Proceedings of the National Academy of Sciences of the United States of America 104, 18265–9.

Luck, S., 2005. An introduction to the event-related potential technique. MIT Press.

Luck, S.J., Kappenman, E.S., 2012. The Oxford Handbook of Event-Related Potential Components. Oxford University Press.

Lynch, J.C., Mclaren, J.W., 1989. Deficits of visual attention and saccadic eye movements after lesions of parietooccipital cortex in monkeys Deficits of Visual Attention and Saccadic Eye Movements After Lesions of Parietooccipital Cortex in Monkeys. Journal of Neurophysiology 61, 74–90.

MacMillan, N.A., Creelman, C.D., 2005. MacMillan, N.A. & Creelman, C.D., 2005. Detection Theory: A User's Guide. Psychology Press.

Makeig, S., Debener, S., Onton, J., Delorme, A., 2004. Mining event-related brain dynamics. Trends in cognitive sciences 8, 204–10.

Maldjian, J.A., Laurienti, P.J., Kraft, R.A., Burdette, J.H., 2003. An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. NeuroImage 19, 1233–9.

Malhotra, P.A., Parton, A.D., Greenwood, R., Husain, M., 2006. Noradrenergic modulation of space exploration in visual neglect. Annals of Neurology 59, 186–90.

Manly, T., Dobler, V.B., Dodds, C.M., George, M.A., 2005. Rightward shift in spatial awareness with declining alertness. Neuropsychologia 43, 1721–8.

Mantini, D., Perrucci, M.G., Del Gratta, C., Romani, G.L., Corbetta, M., 2007. Electrophysiological signatures of resting state networks in the human brain. Proceedings of the National Academy of Sciences of the United States of America 104, 13170–5.

Mantini, D., Vanduffel, W., 2013. Emerging roles of the brain's default network. The Neuroscientist 19, 76–87.

Margulies, D.S., Kelly, A.M.C., Uddin, L.Q., Biswal, B.B., Castellanos, F.X., Milham, M.P., 2007. Mapping the functional connectivity of anterior cingulate cortex. NeuroImage 37, 579–88.

Maris, E., Oostenveld, R., 2007. Nonparametric statistical testing of EEG- and MEG-data. Journal of neuroscience methods 164, 177–90.

Martin, T., Huxlin, K., Kavcic, V., 2010. Motion-onset visual evoked potentials predict performance during a global direction discrimination task. Neuropsychologia 48, 3563–72.

Marx, E., Deutschländer, A., Stephan, T., Dieterich, M., Wiesmann, M., Brandt, T., 2004. Eyes open and eyes closed as rest conditions: impact on brain activation patterns. NeuroImage 21, 1818–24.

Mason, M.F., Norton, M.I., Van Horn, J.D., Wegner, D.M., Grafton, S.T., Macrae, C.N., 2007. Wandering minds: the default network and stimulus-independent thought. Science (New York, N.Y.) 315, 393–5.

Mathews, R., Buss, R., Stanley, W., Fields, B., Cho, J., Bruhan, B., 1989. Role of implicit and explicit processes of learning from examples: A synergistic effect. Journal of Experimental Psychology: Learning, Memory, and Cognition 15, 1083 – 1100.

Mathewson, K.E., Gratton, G., Fabiani, M., Beck, D.M., 2009. To See or Not to See: Pre-stimulus Alpha Phase Predicts Visual Awareness. The Journal of Neuroscience 29, 2725–2732.

Matthias, E., Bublak, P., Costa, A., Müller, H.J., Schneider, W.X., Finke, K., 2009. Attentional and sensory effects of lowered levels of intrinsic alertness. Neuropsychologia 47, 3255–64.

Mattingley, J.B., Pierson, J.M., Bradshaw, J.L., Phillips, J.G., Bradshaw, J.A., 1993. To see or not to see: The effects of visible and invisible cues on line bisection judgements in unilateral neglect. Neuropsychologia 31, 1201–1215.

May, A., 2011. Experience-dependent structural plasticity in the adult human brain. Trends in cognitive sciences 15, 475–82.

Mazaheri, A., Nieuwenhuis, I.L.C., van Dijk, H., Jensen, O., 2009. Prestimulus alpha and mu activity predicts failure to inhibit motor responses. Human brain mapping 30, 1791–800.

McGuire, P.K., Paulesu, E., Frackowiak, R.S., Frith, C.D., 1996. Brain activity during stimulus independent thought. Neuroreport 7, 2095–9.

McIntosh, A.R., Gonzalez-Lima, F., 1991. Structural modeling of functional neural pathways mapped with 2-deoxyglucose: effects of acoustic startle habituation on the auditory system. Brain research 547, 295–302.

McKeown, M.J., Makeig, S., Brown, G.G., Jung, T.P., Kindermann, S.S., Bell, A.J., Sejnowski, T.J., 1998. Analysis of fMRI data by blind separation into independent spatial components. Human brain mapping 6, 160–88.

McManus, I.C., Freegard, M., Moore, J., Rawles, R., 2010. Science in the Making: Right Hand, Left Hand. II: The duck-rabbit figure. Laterality 15, 166–85.

McMillan, R.L., Kaufman, S.B., Singer, J.L., 2013. Ode to positive constructive daydreaming. Frontiers in Psychology 4, 626.

McVay, J.C., Kane, M.J., Kwapil, T.R., 2009. Tracking the train of thought from the laboratory into everyday life: an experience-sampling study of mind wandering across controlled and ecological contexts. Psychonomic bulletin & review 16, 857–63.

Meindl, T., Teipel, S., Elmouden, R., Mueller, S., Koch, W., Dietrich, O., Coates, U., Reiser, M., Glaser, C., 2010. Test-retest reproducibility of the default-mode network in healthy individuals. Human brain mapping 31, 237–46.

Mennes, M., Kelly, C., Zuo, X.-N., Di Martino, A., Biswal, B.B., Castellanos, F.X., Milham, M.P., 2010. Inter-individual differences in resting-state functional connectivity predict task-induced BOLD activity. NeuroImage 50, 1690–701.

Miall, R.C., Robertson, E.M., 2006. Functional imaging: is the resting brain resting? Current biology : CB 16, R998–1000.

Miezin, F.M., Maccotta, L., Ollinger, J.M., Petersen, S.E., Buckner, R.L., 2000. Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. NeuroImage 11, 735–59.

Mintun, M.A., Lundstrom, B.N., Snyder, A.Z., Vlassenko, A.G., Shulman, G.L., Raichle, M.E., 2001. Blood flow and oxygen delivery to human brain during functional activity: theoretical modeling and experimental data. Proceedings of the National Academy of Sciences of the United States of America 98, 6859–64.

Mitra, P.P., Pesaran, B., 1999. Analysis of dynamic brain imaging data. Biophysical journal 76, 691–708.

Mo, J., Schroeder, C.E., Ding, M., 2011. Attentional modulation of alpha oscillations in macaque inferotemporal cortex. The Journal of neuroscience : the official journal of the Society for Neuroscience 31, 878–82.

Monto, S., Palva, S., Voipio, J., Palva, J.M., 2008. Very slow EEG fluctuations predict the dynamics of stimulus detection and oscillation amplitudes in humans. The Journal of neuroscience : the official journal of the Society for Neuroscience 28, 8268–72.

Moosmann, M., Ritter, P., Krastel, I., Brink, A., Thees, S., Blankenburg, F., Taskin, B., Obrig, H., Villringer, A., 2003. Correlates of alpha rhythm in functional magnetic resonance imaging and near infrared spectroscopy. NeuroImage 20, 145–58.

Murphy, K., Birn, R.M., Handwerker, D.A., Jones, T.B., Bandettini, P.A., 2009. The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? NeuroImage 44, 893–905.

Necker, L.A., 1832. Observations on some remarkable optical phaenomena seen in Switzerland; and on an optical phaenomenon which occurs on viewing a figure of a crystal or geometrical solid. London and Edinburgh Philosophical Magazine and Journal of Science 1, 329–337.

Neves, G., Cooke, S.F., Bliss, T.V.P., 2008. Synaptic plasticity, memory and the hippocampus: a neural network approach to causality. Nature reviews. Neuroscience 9, 65–75.

Newsome, W.T., Salzman, C.D., 1993. The neuronal basis of motion perception. Ciba Foundation symposium 174, 217–30; discussion 230–46.

Niedeggen, M., Wist, E.R., 1999. Characteristics of visual evoked potentials generated by motion coherence onset. Brain research. Cognitive brain research 8, 95–105.

Nir, Y., Mukamel, R., Dinstein, I., Privman, E., Harel, M., Fisch, L., Gelbard-Sagiv, H., Kipervasser, S., Andelman, F., Neufeld, M.Y., Kramer, U., Arieli, A., Fried, I., Malach, R., 2008. Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex. Nature neuroscience 11, 1100–8.

Nobre, A., Correa, A., Coull, J., 2007. The hazards of time. Current opinion in neurobiology 17, 465–70.

Olshausen, B.A., 2003. Principles of image representation in visual cortex., in: Chalupa, L.M., Werner, J.S. (Eds.), The Visual Neurosciences. MIT Press, Cambridge, Mass., pp. 1603–15.

Oostenveld, R., Fries, P., Maris, E., Schoffelen, J.-M., 2011. FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. Computational intelligence and neuroscience 2011, 156869.

Oosterhof, N.N., Todorov, A., 2008. The functional basis of face evaluation. Proceedings of the National Academy of Sciences of the United States of America 105, 11087–92.

Packard, M.G., Knowlton, B.J., 2002. Learning and memory functions of the Basal Ganglia. Annual review of neuroscience 25, 563–93.

Penfield, W., Milner, B., 1958. Memory deficit produced by bilateral lesions in the hippocampal zone. A.M.A. archives of neurology and psychiatry 79, 475–97.

Penny, W.D., Stephan, K.E., Mechelli, a, Friston, K.J., 2004. Comparing dynamic causal models. NeuroImage 22, 1157–72.

Piantoni, G., Kline, K.A., Eagleman, D.M., 2010. Beta oscillations correlate with the probability of perceiving rivalrous visual stimuli. Journal of Vision 10, 1–11.

Poldrack, R.A., Packard, M.G., 2003. Competition among multiple memory systems: converging evidence from animal and human brain studies. Neuropsychologia 41, 245–51.

Posner, M., Raichle, M., 1994. Images of Mind. W.H. Freeman and Company, New York.

Posner, M.I., Walker, J.A., Friedrich, F.J., Rafal, R.D., 1984. Effects of parietal injury on covert orienting of attention. Journal of Neuroscience 4, 1863–74.

Pouget, A., Dayan, P., Zemel, R.S., 2003. Inference and computation with population codes. Annual review of neuroscience 26, 381–410.

Rafal, R.D., 1994. Neglect. Current Opinion in Neurobiology 4, 231–6.

Raichle, M.E., 2010. Two views of brain function. Trends in cognitive sciences 14, 180–90.

Raichle, M.E., 2011. The restless brain. Brain connectivity 1, 3–12.

Raichle, M.E., MacLeod, A.M., Snyder, A.Z., Powers, W.J., Gusnard, D.A., Shulman, G.L., 2001. A default mode of brain function. Proceedings of the National Academy of Sciences of the United States of America 98, 676–82.

Raichle, M.E., Mintun, M.A., 2006. Brain work and brain imaging. Annual review of neuroscience 29, 449–76.

Raichle, M.E., Snyder, A.Z., 2007. A default mode of brain function: a brief history of an evolving idea. NeuroImage 37, 1083–90; discussion 1097–9.

Ramot, M., Wilf, M., Goldberg, H., Weiss, T., Deouell, L.Y., Malach, R., 2011. Coupling between spontaneous (resting state) fMRI fluctuations and human oculo-motor activity. NeuroImage 58, 213–225.

Reber, A.S., 1989. Implicit learning and tacit knowledge. Journal of Experimental Psychology: General 118, 219–35.

Rees, G., Friston, K., Koch, C., 2000. A direct quantitative relationship between the functional properties of human and macaque V5. Nature Neuroscience 3, 716–23.

Rees, G., Kreiman, G., Koch, C., 2002. Neural correlates of consciousness in humans. Nature reviews. Neuroscience 3, 261–70.

Rees, G., Wojciulik, E., Clarke, K., Husain, M., Frith, C., Driver, J., 2000. Unconscious activation of visual cortex in the damaged right hemisphere of a parietal patient with extinction. Brain 123 (Pt 8, 1624–33.

Rees, G., Wojciulik, E., Clarke, K., Husain, M., Frith, C., Driver, J., 2002. Neural correlates of conscious and unconscious vision in parietal extinction. Neurocase 8, 387–93.

Reiss, J.P., Campbell, D.W., Leslie, W.D., Paulus, M.P., Stroman, P.W., Polimeni, J.O., Malcolmson, K.A., Sareen, J., 2005. The role of the striatum in implicit learning: a functional magnetic resonance imaging study. Neuroreport 16, 1291–5.

Ress, D., Backus, B.T., Heeger, D.J., 2000. Activity in primary visual cortex predicts performance in a visual detection task. Nature Neuroscience 3, 940–5.

Robertson, I.H., Mattingley, J.B., Rorden, C., Driver, J., 1998. Phasic alerting of neglect patients overcomes their spatial deficit in visual awareness. Nature 395, 169–72.

Rockstroh, B., Elbert, T., Canavan, A.G., Lutzenberger, W., Birbaumer, N., 1989. Slow Cortical Potentials and Behavior,, Second. ed. Urban & Schwarzenberg, Baltimore.

Romei, V., Brodbeck, V., Michel, C., Amedi, A., Pascual-Leone, A., Thut, G., 2008. Spontaneous fluctuations in posterior alpha-band EEG activity reflect variability in excitability of human visual areas. Cerebral cortex (New York, N.Y. : 1991) 18, 2010–8.

Rose, M., Haider, H., Salari, N., Büchel, C., 2011. Functional dissociation of hippocampal mechanism during implicit learning based on the domain of associations. The Journal of Neuroscience 31, 13739–45.

Rubin, E., 1915. Visuell Wahrgenommene Figuren. Gyldenalske Boghandel, Copenhagen.

Sadaghiani, S., Hesselmann, G., Friston, K.J., Kleinschmidt, A., 2010. The relation of ongoing brain activity, evoked neural responses, and cognition. Frontiers in systems neuroscience 4, 20.

Sadaghiani, S., Hesselmann, G., Kleinschmidt, A., 2009. Distributed and antagonistic contributions of ongoing activity fluctuations to auditory stimulus detection. The Journal of neuroscience : the official journal of the Society for Neuroscience 29, 13410–7.

Sadaghiani, S., Scheeringa, R., Lehongre, K., Morillon, B., Giraud, A.-L., Kleinschmidt, A., 2010. Intrinsic connectivity networks, alpha oscillations, and tonic alertness: a simultaneous electroencephalography/functional magnetic resonance imaging study. The Journal of neuroscience : the official journal of the Society for Neuroscience 30, 10243–50.

Saleh, M., Reimer, J., Penn, R., Ojakangas, C.L., Hatsopoulos, N.G., 2010. Fast and slow oscillations in human primary motor cortex predict oncoming behaviorally relevant cues. Neuron 65, 461–71.

Salinas, E., Sejnowski, T.J., 2001. Correlated neuronal activity and the flow of neural information. Nature reviews. Neuroscience 2, 539–50.

Sapir, A., d'Avossa, G., McAvoy, M., Shulman, G.L., Corbetta, M., 2005. Brain signals for spatial attention predict performance in a motion discrimination task.

Proceedings of the National Academy of Sciences of the United States of America 102, 17810–5.

Sarri, M., Kalra, L., Greenwood, R., Driver, J., 2006. Prism adaptation changes perceptual awareness for chimeric visual objects but not for chimeric faces in spatial neglect after right-hemisphere stroke. Neurocase 12, 127–35.

Sarri, M., Ruff, C.C., Rees, G., Driver, J., 2010. Neural correlates of visual extinction or awareness in a series of patients with right temporoparietal damage. Cognitive Neuroscience 1, 16–25.

Schapiro, A.C., Kustner, L. V, Turk-Browne, N.B., 2012. Shaping of object representations in the human medial temporal lobe based on temporal regularities. Current Biology 22, 1622–7.

Schendan, H.E., Searl, M.M., Melrose, R.J., Stern, C.E., 2003. An FMRI study of the role of the medial temporal lobe in implicit and explicit sequence learning. Neuron 37, 1013–25.

Schölvinck, M.L., Friston, K.J., Rees, G., 2012. The influence of spontaneous activity on stimulus processing in primary visual cortex. NeuroImage 59, 2700–8.

Schölvinck, M.L., Leopold, D.A., Brookes, M.J., Khader, P.H., 2013. The contribution of electrophysiology to functional connectivity mapping. NeuroImage 80, 297–306.

Schölvinck, M.L., Maier, A., Ye, F.Q., Duyn, J.H., Leopold, D.A., 2010. Neural basis of global resting-state fMRI activity. Proceedings of the National Academy of Sciences of the United States of America 107, 10238–43.

Scholz, J., Klein, M.C., Behrens, T.E.J., Johansen-Berg, H., 2009. Training induces changes in white-matter architecture. Nature neuroscience 12, 1370–1.

Schooler, J.W., Raichle, M.E., Halpern, D.V., 2004. Thinking and Seeing: Visual metacognition in Adults and Children, in: Levine, D.T. (Ed.), MIT Press, Cambridge, Mass., pp. 203–26.

Schooler, J.W., Smallwood, J., Christoff, K., Handy, T.C., Reichle, E.D., Sayette, M. a, 2011. Meta-awareness, perceptual decoupling and the wandering mind. Trends in cognitive sciences 15, 319–26.

Schroeder, C.E., Lakatos, P., 2009. Low-frequency neuronal oscillations as instruments of sensory selection. Trends in neurosciences 32, 9–18.

Schultz, W., 2007. Multiple dopamine functions at different time courses. Annual review of neuroscience 30, 259–88.

Schwarzkopf, D.S., Song, C., Rees, G., 2011. The surface area of human V1 predicts the subjective experience of object size. Nature neuroscience 14, 28–30.

Seeley, W.W., Menon, V., Schatzberg, A.F., Keller, J., Glover, G.H., Kenna, H., Reiss, A.L., Greicius, M.D., 2007. Dissociable intrinsic connectivity networks for salience processing and executive control. The Journal of neuroscience : the official journal of the Society for Neuroscience 27, 2349–56.

Seitz, A., Watanabe, T., 2005. A unified model for perceptual learning. Trends in cognitive sciences 9, 329–34.

Seitz, A.R., Yamagishi, N., Werner, B., Goda, N., Kawato, M., Watanabe, T., 2005. Task-specific disruption of perceptual learning. Proceedings of the National Academy of Sciences of the United States of America 102, 14895–900.

Senkowski, D., Gomez-Ramirez, M., Lakatos, P., Wylie, G.R., Molholm, S., Schroeder, C.E., Foxe, J.J., 2007. Multisensory processing and oscillatory activity: analyzing non-linear electrophysiological measures in humans and simians. Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale 177, 184–95.

Shehzad, Z., Kelly, A.M.C., Reiss, P.T., Gee, D.G., Gotimer, K., Uddin, L.Q., Lee, S.H., Margulies, D.S., Roy, A.K., Biswal, B.B., Petkova, E., Castellanos, F.X., Milham, M.P., 2009. The resting brain: unconstrained yet reliable. Cerebral cortex (New York, N.Y. : 1991) 19, 2209–29.

Shibata, K., Watanabe, T., Sasaki, Y., Kawato, M., 2011. Perceptual learning incepted by decoded fMRI neurofeedback without stimulus presentation. Science 334, 1413–5.

Shmuel, A., Leopold, D.A., 2008. Neuronal correlates of spontaneous fluctuations in fMRI signals in monkey visual cortex: Implications for functional connectivity at rest. Human brain mapping 29, 751–61.

Shtyrov, Y., 2012. Neural bases of rapid word learning. The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry 18, 312–9.

Shulman, G.L., Fiez, J.A., Corbetta, M., Buckner, R.L., Miezin, F.M., Raichle, M.E., Petersen, S.E., 1997. Common Blood Flow Changes across Visual Tasks: II. Decreases in Cerebral Cortex. Journal of cognitive neuroscience 9, 648–63.

Siegel, M., Donner, T.H., Oostenveld, R., Fries, P., Engel, A.K., 2007. High-frequency activity in human visual cortex is modulated by visual motion strength. Cerebral cortex (New York, N.Y. : 1991) 17, 732–41.

Singer, J.L., Antrobus, J.S., 1972. Daydreaming; imaginal processes; and personality: A normative study, in: Sheehan, P.W. (Ed.), The Function and Nature of Imagery. Academic Press, New York, pp. 175–202.

Sio, U.N., Ormerod, T.C., 2009. Does incubation enhance problem solving? A meta-analytic review. Psychological bulletin 135, 94–120.

Smallwood, J., Brown, K., Baird, B., Schooler, J.W., 2012. Cooperation between the default mode network and the frontal-parietal network in the production of an internal train of thought. Brain research 1428, 60–70.

Smallwood, J., Davies, J.B., Heim, D., Finnigan, F., Sudberry, M., O'Connor, R., Obonsawin, M., 2004. Subjective experience and the attentional lapse: task engagement and disengagement during sustained attention. Consciousness and cognition 13, 657–90.

Smallwood, J., Fishman, D.J., Schooler, J.W., 2007. Counting the cost of an absent mind: mind wandering as an underrecognized influence on educational performance. Psychonomic bulletin & review 14, 230–6.

Smallwood, J., McSpadden, M., Luus, B., Schooler, J., 2008. Segmenting the stream of consciousness: the psychological correlates of temporal structures in the time series data of a continuous performance task. Brain and cognition 66, 50–6.

Smallwood, J., McSpadden, M., Schooler, J.W., 2007. The lights are on but no one's home: meta-awareness and the decoupling of attention when the mind wanders. Psychonomic bulletin & review 14, 527–33.

Smallwood, J., Schooler, J.W., 2006. The restless mind. Psychological bulletin 132, 946–58.

Smallwood, J.M., Baracaia, S.F., Lowe, M., Obonsawin, M., 2003. Task unrelated thought whilst encoding information. Consciousness and Cognition 12, 452–484.

Smallwood, J.M., McSpadden, M., Schooler, J.W., 2008. When attention matters: The curious incident of the wandering mind. Memory & Cognition 36, 1144–50.

Smilek, D., Carriere, J.S. a, Cheyne, J.A., 2010. Out of mind, out of sight: eye blinking as indicator and embodiment of mind wandering. Psychological science 21, 786–9.

Smith, A.P.R., Stephan, K.E., Rugg, M.D., Dolan, R.J., 2006. Task and content modulate amygdala-hippocampal connectivity in emotional retrieval. Neuron 49, 631–8.

Smith, J.E.T., Zhan, C. a, Cook, E.P., 2011. The functional link between area MT neural fluctuations and detection of a brief motion stimulus. The Journal of neuroscience : the official journal of the Society for Neuroscience 31, 13458–68.

Smith, K., 2012. Neuroscience: Idle minds. Nature 489, 356–8.

Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., Beckmann, C.F., 2009a. Correspondence of the brain's functional architecture during activation and rest. Proceedings of the National Academy of Sciences of the United States of America 106, 13040–5.

Smith, S.M., Fox, P.T., Miller, K.L., Glahn, D.C., Fox, P.M., Mackay, C.E., Filippini, N., Watkins, K.E., Toro, R., Laird, A.R., Beckmann, C.F., 2009b. Correspondence of the brain's functional architecture during activation and rest. Proceedings of the National Academy of Sciences of the United States of America 106, 13040–5.

Sporns, O., Tononi, G., Edelman, G.M., 2000. Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices. Cerebral cortex (New York, N.Y. : 1991) 10, 127–41.

Sprague, J.M., 1966. Interaction of cortex and superior colliculus in mediation of visually guided behavior in the cat. Science 153, 1544–1547.

Stawarczyk, D., Majerus, S., Maquet, P., D'Argembeau, A., 2011. Neural correlates of ongoing conscious experience: both task-unrelatedness and stimulus-independence are related to default network activity. PloS one 6, e16997.

Stefanovic, B., Schwindt, W., Hoehn, M., Silva, A.C., 2007. Functional uncoupling of hemodynamic from neuronal response by inhibition of neuronal nitric oxide synthase. Journal of cerebral blood flow and metabolism : official journal of the International Society of Cerebral Blood Flow and Metabolism 27, 741–54.

Stephan, K.E., 2004. On the role of general system theory for functional neuroimaging. Journal of anatomy 205, 443–70.

Stephan, K.E., Friston, K.J., 2011. Analyzing effective connectivity with fMRI. Wiley Interdiscip Rev Cogn Sci 1, 446–459.

Stephan, K.E., Harrison, L.M., Penny, W.D., Friston, K.J., 2004. Biophysical models of fMRI responses. Current opinion in neurobiology 14, 629–35.

Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., Friston, K.J., 2009. Bayesian model selection for group studies. NeuroImage 46, 1004–17.

Stephan, K.E., Weiskopf, N., Drysdale, P.M., Robinson, P.A., Friston, K.J., 2007. Comparing hemodynamic models with DCM. NeuroImage 38, 387–401.

Stevens, W.D., Buckner, R.L., Schacter, D.L., 2010. Correlated low-frequency BOLD fluctuations in the resting human brain are modulated by recent

experience in category-preferential visual regions. Cerebral Cortex 20, 1997–2006.

Summerfield, C., Egner, T., Greene, M., Koechlin, E., Mangels, J., Hirsch, J., 2006. Predictive codes for forthcoming perception in the frontal cortex. Science 314, 1311–4.

Sun, R., Thomas, M., McClelland, J., Griffiths, T., Kemp, C., Tenenbaum, J., Schoener, G., Bringsjord, S., Taatgen, N., Anderson, J., Norman, K., Detre, G., Polyn, S., Rogers, T., Kruschke, J., Busemeyer, J., Johnson, J., Heit, E., Johnson-Laird, P., Yang, Y., Cleeremans, A., Dienes, Z., Pisapia, N. De, Repovs, G., Braver, T., Shultz, T., Sirois, S., Chater, N., Christiansen, M., Read, S., Monroe, B., Thagard, P., Litt, A., Gray, W., Lopez, F., Shanks, D., Sinha, P., Balas, B., Mussa-Ivaldi, F., Solla, S., 2008. The Cambridge Handbook of Computational Psychology. Cambridge University Press, Cambridge.

Supèr, H., van der Togt, C., Spekreijse, H., Lamme, V.A.F., 2003. Internal state of monkey primary visual cortex (V1) predicts figure-ground perception. The Journal of neuroscience : the official journal of the Society for Neuroscience 23, 3407–14.

Sutton, R., Barto, A., 1998. Reinforcement learning: an introduction. MIT Press, Cambridge, Mass.

Tadin, D., Silvanto, J., Pascual-Leone, A., Battelli, L., 2011. Improved motion perception and impaired spatial suppression following disruption of cortical area MT/V5. Journal of Neuroscience 31, 1279–1283.

Takano, T., Tian, G.-F., Peng, W., Lou, N., Libionka, W., Han, X., Nedergaard, M., 2006. Astrocyte-mediated control of cerebral blood flow. Nature neuroscience 9, 260–7.

Tambini, A., Ketz, N., Davachi, L., 2010. Enhanced brain correlations during rest are related to memory for recent experiences. Neuron 65, 280–90.

Taubert, M., Lohmann, G., Margulies, D.S., Villringer, A., Ragert, P., 2011. Long-term effects of motor training on resting-state networks and underlying brain structure. NeuroImage 57, 1492–8.

Taylor, D.G., Bushell, M.C., 1985. The spatial mapping of translational diffusion coefficients by the NMR imaging technique. Physics in medicine and biology 30, 345–9.

Thut, G., Nietzel, A., Brandt, S. a, Pascual-Leone, A., 2006. Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. The Journal of neuroscience : the official journal of the Society for Neuroscience 26, 9494–502.

Thut, G., Veniero, D., Romei, V., Miniussi, C., Schyns, P., Gross, J., 2011. Rhythmic TMS causes local entrainment of natural oscillatory signatures. Current biology : CB 21, 1176–85.

Tohka, J., Foerde, K., Aron, A.R., Tom, S.M., Toga, A.W., Poldrack, R.A., 2008. Automatic independent component labeling for artifact removal in fMRI. NeuroImage 39, 1227–45.

Tomassini, V., Jbabdi, S., Kincses, Z.T., Bosnell, R., Douaud, G., Pozzilli, C., Matthews, P.M., Johansen-Berg, H., 2011. Structural and functional bases for individual differences in motor learning. Human brain mapping 32, 494–508.

Torrance, E.P., 1962. Guiding creative talent. Prentice Hall, Englewood Cliffs, NJ.

Tsodyks, M., Kenet, T., Grinvald, A., Arieli, A., 1999. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. Science (New York, N.Y.) 286, 1943–6.

Turk-Browne, N.B., Scholl, B.J., Johnson, M.K., Chun, M.M., 2010. Implicit perceptual anticipation triggered by statistical learning. The Journal of Neuroscience 30, 11177–87.

Uhlhaas, P.J., Roux, F., Rodriguez, E., Rotarska-Jagiela, A., Singer, W., 2010. Neural synchrony and the development of cortical networks. Trends in Cognitive Sciences 14, 72–80.

Umarova, R.M., Saur, D., Kaller, C.P., Vry, M.-S., Glauche, V., Mader, I., Hennig, J., Weiller, C., 2011. Acute visual neglect and extinction: distinct functional state of the visuospatial attention system. Brain 134, 3310–25.

Vaden, R.J., Hutcheson, N.L., McCollum, L.A., Kentros, J., Visscher, K.M., 2012. Older adults, unlike younger adults, do not modulate alpha power to suppress irrelevant information. NeuroImage 63, 1127–33.

Vahdat, S., Darainy, M., Milner, T.E., Ostry, D.J., 2011. Functionally specific changes in resting-state sensorimotor networks after motor learning. The Journal of neuroscience : the official journal of the Society for Neuroscience 31, 16907–15.

Vaina, L.M., Belliveau, J.W., des Roziers, E.B., Zeffiro, T. a, 1998. Neural systems underlying learning and representation of global motion. Proceedings of the National Academy of Sciences of the United States of America 95, 12657–62.

Vaina, L.M., Cowey, A., Eskew, R.T., LeMay, M., Kemper, T., 2001. Regional cerebral correlates of global motion perception: evidence from unilateral cerebral brain damage. Brain 124, 310–21.

Vaina, L.M., Sundareswaran, V., Harris, J.G., 1995. Learning to ignore: psychophysics and computational modeling of fast learning of direction in noisy motion stimuli. Brain Research Cognitive Brain Research 2, 155–63.

Vallar, G., 1993. The anatomical basis of spatial hemineglect in humans., in: Robertson, I.H., Marshall, J.C. (Eds.), Unilateral Neglect: Clinical and Experimental Studies. Lawrence Erlbaum, Hillsdale (NJ), pp. 27–59.

Vallar, G., Rusconi, M.L., Bignamini, L., Geminiani, G., Perani, D., 1994. Anatomical correlates of visual and tactile extinction in humans: a clinical CT scan study. Journal of Neurology, Neurosurgery, and Psychiatry 57, 464–70.

Van de Ven, V.G., Formisano, E., Prvulovic, D., Roeder, C.H., Linden, D.E.J., 2004. Functional connectivity as revealed by spatial independent component analysis of fMRI measurements during rest. Human brain mapping 22, 165–78.

Van Dijk, H., Schoffelen, J.-M., Oostenveld, R., Jensen, O., 2008. Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability. The Journal of neuroscience : the official journal of the Society for Neuroscience 28, 1816–23.

Van Dijk, K.R.A., Hedden, T., Venkataraman, A., Evans, K.C., Lazar, S.W., Buckner, R.L., 2010. Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. Journal of neurophysiology 103, 297–321.

Van Ede, F., Jensen, O., Maris, E., 2010. Tactile expectation modulates pre-stimulus beta-band oscillations in human sensorimotor cortex. NeuroImage 51, 867–76.

Van Opstal, F., Verguts, T., Orban, G.A., Fias, W., 2008. A hippocampal-parietal network for learning an ordered sequence. NeuroImage 40, 333–41.

Vanhaudenhuyse, A., Noirhomme, Q., Tshibanda, L.J.-F., Bruno, M.-A., Boveroux, P., Schnakers, C., Soddu, A., Perlbarg, V., Ledoux, D., Brichant, J.-F., Moonen, G., Maquet, P., Greicius, M.D., Laureys, S., Boly, M., 2010. Default network connectivity reflects the level of consciousness in non-communicative brain-damaged patients. Brain : a journal of neurology 133, 161–71.

Vincent, J.L., Patel, G.H., Fox, M.D., Snyder, A.Z., Baker, J.T., Van Essen, D.C., Zempel, J.M., Snyder, L.H., Corbetta, M., Raichle, M.E., 2007. Intrinsic functional architecture in the anaesthetized monkey brain. Nature 447, 83–86.

Vogels, R., Spileers, W., Orban, G.A., 1989. The response variability of striate cortical neurons in the behaving monkey. Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale 77, 432–6.

Vogt, B.A., Vogt, L., Laureys, S., 2006. Cytology and functionally correlated circuits of human posterior cingulate areas. NeuroImage 29, 452–66.

Vossel, S., Eschenbeck, P., Weiss, P.H., Weidner, R., Saliger, J., Karbe, H., Fink, G.R., 2011. Visual extinction in relation to visuospatial neglect after right-hemispheric stroke: quantitative assessment and statistical lesion-symptom mapping. Journal of Neurology, Neurosurgery, and Psychiatry 82, 862–8.

Vuilleumier, P., Armony, J.L., Clarke, K., Husain, M., Driver, J., Dolan, R.J., 2002. Neural response to emotional faces with and without awareness: event-related fMRI in a parietal patient with visual extinction and spatial neglect. Neuropsychologia 40, 2156–66.

Vuilleumier, P., Schwartz, S., Verdon, V., Maravita, A., Husain, M., Driver, J., 2010. Attention-dependent functional abnormality in retinotopic visual cortex for patients with parietal lesions and spatial neglect. Current Biology 18, 1525–1529.

Vuilleumier, P.O., Rafal, R.D., 2000. A systematic study of visual extinction. Between- and within-field deficits of attention in hemispatial neglect. Brain 123 (Pt 6, 1263–79.

Vul, E., Pashler, H., 2007. Incubation benefits only after people have been misdirected. Memory & cognition 35, 701–10.

Waites, A.B., Stanislavsky, A., Abbott, D.F., Jackson, G.D., 2005. Effect of prior cognitive state on resting state networks measured with functional connectivity. Human brain mapping 24, 59–68.

Walsh, V., Ellison, A., Ashbridge, E., Cowey, A., 1999. The role of the parietal cortex in visual attention--hemispheric asymmetries and the effects of learning: a magnetic stimulation study. Neuropsychologia 37, 245–51.

Wang, G., Grone, B., Colas, D., Appelbaum, L., Mourrain, P., 2011. Synaptic plasticity in sleep: learning, homeostasis and disease. Trends in neurosciences 34, 452–63.

Wang, J., Zuo, X., He, Y., 2010. Graph-based network analysis of resting-state functional MRI. Frontiers in systems neuroscience 4, 16.

Wang, X.-J., 2010. Neurophysiological and computational principles of cortical rhythms in cognition. Physiological reviews 90, 1195–268.

Wenger, K.K., Visscher, K.M., Miezin, F.M., Petersen, S.E., Schlaggar, B.L., 2004. Comparison of sustained and transient activity in children and adults using a mixed blocked/event-related fMRI design. NeuroImage 22, 975–85.

White, H.A., Shah, P., 2006. Uninhibited imaginations: Creativity in adults with Attention-Deficit/Hyperactivity Disorder. Personality and Individual Differences 40, 1121–1131.

Wig, G.S., Grafton, S.T., Demos, K.E., Wolford, G.L., Petersen, S.E., Kelley, W.M., 2008. Medial temporal lobe BOLD activity at rest predicts individual differences in memory ability in healthy young adults. Proceedings of the National Academy of Sciences of the United States of America 105, 18555–60.

Wilkinson, L., Jahanshahi, M., 2007. The striatum and probabilistic implicit sequence learning. Brain research 1137, 117–30.

Wilson, B., Cockburn, J., Haligan, P., 1987. Behavioural inattention test. Thames Valley Company, Titchfield, Hants.

Winocur, G., 1985. The hippocampus and thalamus: their roles in short- and long-term memory and the effects of interference. Behavioural brain research 16, 135–52.

Wise, R.G., Ide, K., Poulin, M.J., Tracey, I., 2004. Resting fluctuations in arterial carbon dioxide induce significant low frequency variations in BOLD signal. NeuroImage 21, 1652–64.

Wyart, V., Tallon-Baudry, C., 2009. How ongoing fluctuations in human visual cortex predict perceptual awareness: baseline shift versus decision bias. The Journal of Neuroscience 29, 8715–25.

Yang, A.C., Tsai, S.-J., 2013. Is mental illness complex? From behavior to brain. Progress in neuro-psychopharmacology & biological psychiatry 45, 253–7.

Zang, Y., Jiang, T., Lu, Y., He, Y., Tian, L., 2004. Regional homogeneity approach to fMRI data analysis. NeuroImage 22, 394–400.

Zang, Y.-F., He, Y., Zhu, C.-Z., Cao, Q.-J., Sui, M.-Q., Liang, M., Tian, L.-X., Jiang, T.-Z., Wang, Y.-F., 2007. Altered baseline brain activity in children with ADHD revealed by resting-state functional MRI. Brain & development 29, 83–91.

Zarahn, E., Aguirre, G.K., D'Esposito, M., 1997. Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. NeuroImage 5, 179–97.

Zhang, D., Raichle, M.E., 2010. Disease and the brain's dark energy. Nature reviews. Neurology 6, 15–28.

Zhang, D., Snyder, A.Z., Fox, M.D., Sansbury, M.W., Shimony, J.S., Raichle, M.E., 2008. Intrinsic functional relations between human cerebral cortex and thalamus. Journal of neurophysiology 100, 1740–8.

Zhang, Y., Ding, M., 2010. Detection of a weak somatosensory stimulus: role of the prestimulus mu rhythm and its top-down modulation. Journal of cognitive neuroscience 22, 307–22.

Zhiyan, T., Singer, J.L., 1997. Daydreaming styles, emotionality and the Big Five personality dimensions. Imagination, Cognition and Personality 16, 399–414.

Zohary, E., Shadlen, M.N., Newsome, W.T., 1994. Correlated neuronal discharge rate and its implications for psychophysical performance. Nature 370, 140–3.

Zuo, X.-N., Kelly, C., Adelstein, J.S., Klein, D.F., Castellanos, F.X., Milham, M.P., 2010. Reliable intrinsic connectivity networks: test-retest evaluation using ICA and dual regression approach. NeuroImage 49, 2163–77.