Brain Network Connectivity in Anaesthesia and Disorders of Consciousness

Michael Murphy Craig

Sidney Sussex College

Department of Clinical Neurosciences

University of Cambridge

December 2018

This dissertation is submitted to the University of Cambridge for the degree of Doctor of Philosophy

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text.

It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text.

This thesis does not exceed the 60 000 word limit (including footnotes, references, and appendices) for the Clinical Medicine and Clinical Veterinary Medicine Degree Committee.

The studies outlined in Chapter 2 and Chapter 5 have been published in a peer-reviewed journal and as a peer-reviewed conference paper, respectively. The references for these publications are as follows:

Craig, M.M., Manktelow, A.E., Sahakian, B.J., Menon, D.K., Stamatakis, E.A. (2018). Spectral Diversity in Default Mode Network Connectivity Reflects Behavioral State. *Journal of Cognitive Neuroscience*. 30(4): 526-539

Craig, M.M., Adapa, R., Pappas, I., Menon, D.K., Stamatakis, E.A. (2018). Deep Graph Convolutional Neural Networks Identify Frontoparietal Control and Default Mode Network Contributions to Mental Imagery. *Conference on Computational Cognitive Neuroscience.*

Brain Network Connectivity in Anaesthesia and Disorders of Consciousness

Michael Murphy Craig

Until recently, understanding the nature of consciousness was considered a philosophical pursuit. However, technological developments in brain imaging have allowed the study of consciousness as a natural, neurobiological phenomenon. The neurobiology of consciousness has been studied using cognitive and behavioural testing in healthy volunteers and by examining how brain function and connectivity is altered in various clinical settings. The focus of this thesis is to use two of these clinical settings, pharmacologically-induced sedation and disorders of consciousness (DOC), as experimental models for measuring changes in connectivity patterns associated with alterations in consciousness. Experiment 1 presents a method for improving functional magnetic resonance imaging (fMRI) data pre-processing to measure brain network connectivity more accurately. This pre-processing method is then applied to the analyses in the remainder of the thesis. Experiment 2 focuses on a fMRI dataset in which healthy volunteers were administered propofol, an anaesthetic drug known to act on inhibitory GABAergic interneurons. Using a novel multimodal analysis, changes in functional brain network connectivity in default mode, salience, and frontoparietal control networks were found to correlate with the cortical distribution of parvalbumin-expressing GABAergic interneurons. Using the same dataset, Experiment 3 identified a relationship between structural and functional networks in connections between default mode and salience networks. Similar results have been reported in non-human primate models, however, this is the first study to find network-specific structure-function relationships during sedation in humans. These findings informed the remainder of the thesis, which focused on developing network-based machine learning methods for examining brain connectivity in patients with DOC. Experiment 4 developed and validated a graph convolutional neural network (GCNN) classifier using fMRI data and functional connectivity from healthy volunteers performing a volitional mental imagery task. Experiment 5 applied the GCNN to patients with DOC and found frontoparietal control network connectivity measured at rest to be most important in classifying patients capable of performing the mental imagery task. Taken together, these results contribute to the improvement of brain network analysis techniques, the understanding of the

neurobiology of propofol-induced sedation, and the development of machine learning algorithms to identify DOC patients with preserved covert volitional capacity. This work demonstrates the utility of clinical models in deepening our understanding of the neurobiology of consciousness.

Table of Contents

List of Figures and Tables

Acknowledgements

I would first like to thank my supervisor, Dr Emmanuel Stamatakis, for his encouragement, guidance, and advice over the past few years. His mentorship and intellectual insight have helped me develop into a confident and independent researcher. I am also grateful for Professor David Menon for his support and scientific guidance. My sincerest thanks and appreciation go to Ioannis Pappas for his invaluable help and friendship. I have also benefitted from many thoughtful conversations and insights from past and present members of the Division of Anaesthesia.

I am thankful to Dr Bratislav Mišić for his support during my research visit to the Montreal Neurological Institute. I am also indebted to Dr Dusica Bajic for her mentorship during my time in Boston. I would also like to thank my undergraduate supervisors, Dr Todd Williams, Dr Josef Kurtz, and Dr Joel Kowit, for their encouragement and for giving me the confidence that becoming a scientist was possible.

A special thanks go to many friends who have helped me along the way. I would especially like to thank Toni and Emily for their kindness over the years. I would also like to thank Tim for his friendship and many scientific conversations.

Most of all, though, I am grateful for my family for encouraging and supporting me through the years. To my parents, whose endless love and support have shaped who I am today. To Connor and Jill, for their warmth and dependability when I need it most. To my grandparents, for their kindness and inspiration.

My deepest gratitude goes to my fiancée, Lydia, for her love, encouragement, and brilliance, without which, none of this would be possible. I love you more each day and look forward to seeing what lies ahead for us.

Chapter 1. Review of the Literature

1.1. Introduction

To be conscious is to have subjective experience (Koch, Massimini, Boly, & Tononi, 2016). The overwhelming challenge of defining consciousness has been attempted by many philosophers and scientists over the centuries. Even today, Cartesian dualist accounts are still widely held by the general population and some researchers, wherein conscious states are entirely ineffable and not amenable to physical explanation (Dehaene & Naccache, 2001). Despite this, both theoretical work and empirical research continue, with some consensus around the hypothesis that understanding consciousness represents not a single, but rather multiple problems. Soft problems of consciousness are related to understanding the biophysical correlates of various subjective conscious states, whereas hard problems are concerned with explaining the nature of raw experience (Chalmers, 1995). Despite the compelling distinction drawn from this hypothesis, hard problems of consciousness are, at least with today's technology, impossible to measure. Therefore, this thesis focuses on soft problems of consciousness that are measurable with the tools of modern neuroscience. This is not to imply a reductionist view of consciousness; instead, this work aims to account for the multiple levels of complexity in which the nervous system is organised, with the goal of characterising the neurobiological processes underpinning conscious experience that have implications for clinical medicine. This approach stratifies the conscious state of patients with disorders of consciousness (DOC) and individuals undergoing anaesthesia, providing a model for understanding how these state-changes correlate with underlying brain activity. The following section provides an overview of how consciousness is defined in different clinical settings. Next, the literature surrounding large-scale brain networks is reviewed, as this is the primary neurobiological measure used to characterise these clinical states in this thesis. Finally, an overview of the current state of knowledge of the neuroscience of consciousness is presented.

1.2. Consciousness in the Clinic

Research of the neurobiology of consciousness often uses clinical models to measure changes in brain structure or function that correlate with reduced levels of awareness and wakefulness. This work is not only crucial for understanding consciousness, but also for understanding brain function more generally, with the ultimate goal of translating this work to improve patient care. The first clinical model in this thesis uses an anaesthetic drug to sedate healthy participants. The second clinical model examines patients who have suffered brain injury and are in a state of DOC. The following sections provide context for each of these models, beginning with general anaesthesia.

1.2.1. Overview of General Anaesthesia

General anaesthesia is crucial for safely and humanely performing millions of surgeries each year (Alkire, Hudetz, & Tononi, 2008; Brown, Purdon, & Van Dort, 2011; Rudolph & Antkowiak, 2004). It is characterised by a state of unconsciousness, analgesia, amnesia, and immobility, all while maintaining physiological stability (Akeju & Brown, 2017). Despite its widespread use and significant research in the area, the precise mechanism by which anaesthetic drugs create a state of unconsciousness sufficient to perform surgery is not currently understood (Miller, 2005).

Anaesthesia was first publicly demonstrated in 1846 at Massachusetts General Hospital in Boston using ether (Brown, Purdon, & Van Dort, 2011). Decades of significant technological development have resulted in a medical procedure that can systematically render patients unconscious for an extended period, after which they regain awareness without memory of the procedure. Modern surgery units use a variety of anaesthetic drugs to produce these states. Anaesthetic drugs can be divided into five distinct categories, each of which targets a specific neurotransmitter receptor: gamma-aminobutyric acid type A ($GABA_A$) receptor agonists, *N*-methyl D-aspartate receptor (NMDA) antagonists, α2 adrenergic agonists, dopamine receptor antagonists, and opioid receptor agonists. As each drug type targets a different neurotransmitter system, effects on a patient's physiology, behaviour and subjective

conscious experience vary, resulting in specific drugs being used for certain medical procedures (Brown et al., 2011).

Opioid receptor agonists target μ, κ, and δ opioid receptors in the periaqueductal grey (PAG), rostral ventral medulla (RVM), basal ganglia, amygdala, and spinal cord. Activation of receptors in these regions blocks pain signals moving from the periphery to the brain via the ascending nociceptive pathway (Millan, 2002). Clinically, opioid receptor agonists are primarily used for pain relief during post-operative care or to aid in the maintenance of unconsciousness during anaesthesia (Dahan, Aarts, & Smith, 2010).

NMDA receptor antagonists, which include the drug ketamine, are often used as analgesics or hypnotics in a clinical setting (Bergman, 1999). Ketamine is commonly used in conjunction with other drugs during general anaesthesia or as an alternative to opioids for pain management (Brown et al., 2011). It can also induce dissociative or hallucinatory sensations (Seamans, 2008) and is often reported as a drug of abuse (Morgan, Muetzelfeldt, & Curran, 2009). Previous work suggests that ketamine preferentially binds to NMDA receptors on GABAergic interneurons (Seamans, 2008). This leads to down-regulation of GABAergic inhibition and therefore excitation of pyramidal neurons, resulting in excitatory activity across multiple brain regions (Olney, Newcomer, & Farber, 1999). This process is thought to underlie the dissociative conscious effects of ketamine administration.

Clonidine and dexmedetomidine are examples of α2 adrenergic receptor antagonists (Coursin, Coursin, & Maccioli, 2001), which produce a physiological state similar to non-REM sleep (Brown et al., 2011). Previous studies have shown that the primary target for dexmedetomidine are α2 receptors in the locus coeruleus (Chiu, Chen, Yang, Yang, & Tang, 1995). Dexmedetomidine is used for mechanically ventilating patients in the intensive care unit to provide sedation for relatively short periods (Gerlach & Dasta, 2007). It can also be used for non-intubated patients for various medical and surgical procedures or as an adjunct during general anaesthesia (Carollo, Nossaman, & Ramadhyani, 2008). The subjective experience of sedation using dexmedetomidine is significantly different than that of $GABA_A$ receptor

agonists because patients are easily aroused and have little respiratory depression (Hsu et al., 2004), hence it is thought of as a type of pharmacological non-REM sleep.

Dopamine receptor antagonists, including butyrophenones, haloperidol, and droperidol, can be used as sedatives in conjunction with other anaesthetics and antiemetics. They act on the brain's three main dopaminergic pathways, which include the nigrostriatal pathway, the mesolimbic pathway and the mesocortical pathway (Brown et al., 2011). The alteration in subjective conscious experience as function of dopaminergic receptor antagonist administration is not sufficient for anaesthetic use alone. These drugs are therefore often used in addition to other anaesthetic drugs.

Finally, the primary anaesthetic drugs that are used to induce sedation and unconsciousness during surgical procedures are those that bind to $GABA_A$ receptors (Brown et al., 2011). These GABAergic agonists include sodium thiopental, methohexital, etomidate, and propofol, which act by enhancing the activity of inhibitory interneurons (Rudolph & Antkowiak, 2004). Patients receiving these drugs can become disinhibited, where their behaviour is characterised by euphoria or dysphonia, unintelligible speech, and abnormal movements (Brown, Lydic, & Schiff, 2010). It is typical for patients to enter these states when given low doses of $GABA_A$ agonists during procedures like colonoscopies, dental work, and various radiological procedures (Tung et al., 2001).

Methohexital is rarely used for sedation of patients in the clinic, as small doses have been shown to produce seizures (Metriyakool, 1981). Instead, this effect has led to this drug being used to cortically map seizure activity in patients with epilepsy (Kirchberger et al., 1998). Additionally, methohexital has been used during electroconvulsive therapy (Avramov, White, & Husain, 1995; Vande Voort, Swintak, Wall, & Rasmussen, 2013).

The remaining $GABA_A$ receptor agonists, like propofol and etomidate, are delivered as a bolus in surgical settings, with patients typically losing consciousness within 20-30 seconds. This state is characterised by the patient not being responsive to verbal commands and irregular breathing. Loss of consciousness is often assessed by

having patients count or by having them perform a smooth eye pursuit task. With the loss of consciousness, smooth pursuit decreases and stops, blinking increases, and nystagmus often occurs, though pupillary response to light remains (Brown et al., 2010). Propofol administration results in a concomitant loss of muscle tone, while loss of muscle tone during etomidate administration is preceded by spasms known as myoclonus (Van Keulen & Burton, 2003). Blood pressure also decreases due to the vasodilation and myocardial depression. An anaesthesiologist will then administer a muscle relaxant and intubate the patient. From here, general anaesthesia is maintained with opioids, other hypnotic agents, muscle relaxants, inhalational agents and cardiovascular medications, along with thermoregulatory support and pulmonary ventilation (Brown et al., 2011).

Chapters 3 and 4 in this thesis will focus specifically on propofol-induced sedation. This is because it is one of the most commonly used anaesthetic drugs, produces predictable and stable effects, and is often used as a model to study sedation and consciousness in humans and animals models.

1.2.2. Disorders of Consciousness

DOC are medical conditions, often caused by serious brain injury, that result in inhibition or loss of consciousness (Laureys, Owen, & Schiff, 2004). Focal lesions from traumatic brain injury (TBI) or stroke are a common cause of these disorders, but global injury like anoxia or progressive brain damage from other neurological disorders can also result in a DOC diagnosis. Patients who survive these injuries often begin in a state of unarousable unresponsiveness referred to as a coma. After days or weeks in a coma, patients can recover or succumb to brain death. If neither of these occurs, patients often open their eyes are in one of several different clinically defined states of consciousness (Laureys, Perrin, Schnakers, Boly, & Majerus, 2005). In descending levels of awareness and wakefulness, these disorders include Locked-In Syndrome (LIS), where patients are fully aware and wakeful but incapable of producing motor output (other than eye movements), minimally conscious state (MCS), where patients show some levels of self-awareness on a reproducible and sustained basis, and Unresponsive Wakefulness Syndrome (UWS; also known as

Vegetative State), where patients show signs of wakefulness, but no signs of self-awareness (Laureys et al., 2010, 2004).

In the clinic, neurologists and rehabilitation specialists are tasked with differentiating between these levels of consciousness in unresponsive patients. Due to the severity of the injury and variance of behavioural responses, inconsistent and inaccurate diagnoses often occur, sometimes resulting in inefficient care and incorrect prognosis (Giacino, Fins, Laureys, & Schiff, 2014). Estimating the degree of misdiagnosis is very challenging, with estimations for UWS patients ranging anywhere from 15-45% (Andrews, Murphy, Munday, & Littlewood, 1996; Childs, Mercer, & Childs, 1993; Tresch, Sims, Duthie, Goldstein, & Lane, 1991).

The primary methods for diagnosing patients are based on behavioural assessments. These include methods such as the Glasgow Coma Scale (GCS; Teasdale et al., 2014; Teasdale & Jennet, 1974) and the Consciousness Recovery Scale-Revised score (CRS-R; Giacino et al., 2014; Giacino, Kalmar, & Whyte, 2004; Giacino & Kalmar, 2011), each of which measures a patient's responses to commands. Recent years have seen development of neuroimaging-based methods to supplement behavioural assessments with the goal of improving diagnostic accuracy (Cruse et al., 2011; Di Perri et al., 2016; Giacino et al., 2014; Hannawi, Lindquist, Caffo, Sair, & Stevens, 2015; Laureys et al., 2004; Naci et al., 2014). Clinical care of DOC patients is typically centred on improving various aspects of cognitive capacity and reducing the occurrence of secondary medical complications (Giacino et al., 2014). Improving the diagnostic procedure through additional tests and gaining a deeper understanding of the neurobiology underlying DOC is critical for aiding in patient recovery and developing novel rehabilitation methods. In addition, progress in this field will guide thinking in the various ethical and policy considerations bound to clinical care of patients, such as life-saving treatment and the right-to-die movement (Fins & Shapiro, 2007). The following section outlines the epidemiology, diagnostic procedures, clinical care, and prognosis of patients with DOC.

1.2.2.1. Aetiology and Epidemiology

The most common cause of DOC is TBI (Bernat, 2006). Various types of non-traumatic neuropathology, including hypoxic-ischaemic neuronal injury resulting from cardiac arrest, stroke, and meningitis, can also result in DOC (Tresch et al., 1991). In rare cases, neurodegenerative disorders including Alzheimer's Disease, Huntington's Disease, and Parkinson's Disease can cause DOC (Bernat, 2006; Walshe & Leonard, 1985). Causes of paediatric DOC include perinatal injury, congenital malformations, asphyxia, TBI, and meningitis (Ashwal, Eyman, & Call, 1994).

The prevalence of different types of DOC have been calculated in many studies, however, accurate quantification has proven challenging (Beaumont & Kenealy, 2005). Several published surveys have suggested that in the United States between 56-140 patients per million people suffer from UWS (Bernat, 2006). This is likely a high estimation, given that these surveys included many patients with neurodegenerative disease that were debatable cases of UWS. Further research in this area shows that diagnosis of UWS for patients with neurodegenerative diseases is very inconsistent and that patients who progress to UWS decease shortly thereafter (Bernat, 2006; Volicer, Berman, Cipolloni, & Mandell, 1997). Additionally, these initial surveys included many children with developmental malformations, but not necessarily chronic UWS (Shewmon, Holmes, & Byrne, 1999). A more recent and carefully designed study based on individual assessment of each patient case found the prevalence to be significantly lower, determining that 19 people per million have DOC in Vienna, Austria (Stepan, Haidinger, & Binder, 2004). A cross-sectional study performed between 2000 and 2003 in the Netherlands found an even lower prevalence with only 2 per million people having UWS (Lavrijsen, van den Bosch, Koopmans, & van Weel, 2005). It is not entirely clear why this study found a lower number of UWS cases. One possibility is that different decisions about treatment withdrawal may be more common where the study was performed. Another is that, as this study was performed in nursing homes, more home care of UWS patients could be prominent in the location where the study was done. A more recent study systematically reviewed the literature in April of 2013 and initially examined a total of 1032 studies, with 14 meeting final inclusion criteria (van Erp, Lavrijsen, van de Laar, Vos, & Laureys,

2014). These inclusion criteria consisted of studies that were cross-sectional prevalence studies and included measurements corresponding to the total number of UWS patients in a population. Additionally, they excluded studies that examined specific patient populations (e.g. subarachnoid haemorrhage, cardiac arrest, neurodegenerative disorders), and studies dated before 1994. The authors report that the prevalence of UWS found in these studies was between 2-61 patients per 100,000 people in the population.

The prevalence of MCS is less well studied than UWS. Lomardi et al. (2002) conducted a systematic review of sensory stimulation clinical trials in patients in a coma and UWS (Lombardi, Taricco, De Tanti, Telaro, & Liberati, 2002). Due to the inconsistency in reporting methodological issues, they concluded that accurate assessment of these UWS patients in the trials were poor, and that MCS was likely ten-fold greater than UWS. MCS patient assessment is challenging, so TBI data can be used as a supplement to allow for an estimation of MCS prevalence. There is an incidence of 250 per million moderate-to-severe TBI cases yearly (Bernat, 2006). Of TBI patients with severe injury and GCS scores of 8 or less, 6-16% of patients remained in UWS for 1 month, with approximately 1% remaining in UWS after a year (Braakman, Jennett, & Minderhoud, 1988). Around 10-15% of these patients remained in MCS (Bernat, 2006). Overall, research on understanding the prevalence of DOC is quite limited due to various methodological flaws in their execution, issues surrounding inclusion/exclusion criteria for patient groups and the fact that MCS only became a distinct DOC classification in 2002 (van Erp et al., 2014).

1.2.2.2. Diagnostic and Prognostic Protocols

Diagnosis of DOC can only be made after assessment of a patient's arousal level based on their responses to various behavioural examination tools (Bernat, 2006; Giacino et al., 2014; Koch et al., 2016). One of the earliest tools developed is the GCS (Teasdale et al., 2014; Teasdale & Jennet, 1974). The GCS has been applied extensively to both traumatic (Rowley & Fielding, 1991) and non-traumatic (Mullie et al., 1988) causes of coma and consists of several subscales that measure different behavioural patterns, including eye-opening, verbal responses, and motor responses. The components of each subscale are assessed in a standardised approach that allows for objectivity in

documenting information regarding the patient's level of consciousness. Low scores indicate poor performance suggesting the patient is operating with a reduced level of consciousness (Teasdale et al., 2014). Despite its utility in determining patients in a coma, the GCS is thought to be insufficient for assessing UWS and MCS patients due to its coarse measure of awareness and lack of extensive assessment of neurological functions (Bernat, 2006). More careful evaluation of patient behaviour is critical as it is essential to distinguish between volitional movement that requires awareness and purely reflexive responses to stimuli.

Since the invention of the GCS, a number of tools have been developed to help assess patients with DOC, each with their advantages and disadvantages. The Western neurosensory stimulation profile (Ansell & Keenan, 1989) is sufficient for evaluating patients slowly recovering from TBI, but is less sensitive to differences between patients with lower levels of consciousness such as MCS, UWS, and coma (O'Dell, Jasin, Lyons, Stivers, & Meszaros, 1996). The Wessex Head Injury Scale also aims to help monitor recovery from TBI, with the added benefits of being easy to implement, allowing for collection of data by multidisciplinary rehabilitation teams (Majerus, Vanr De Linden, & Shiel, 2000; Shiel et al., 2000), but is still limited for use with DOC generally.

Additional behavioural assessment tools have been developed specifically for patients with UWS and MCS, however extensive training is required by the medical practitioner applying them. These include the Disorders of Consciousness Scale (DOCS), which measures graded responses to stimuli, as opposed to binary ratings of whether the patient performed a behaviour or not (Pape, Heinemann, Kelly, Hurder, & Lundgren, 2005). It also provides a protocol for mapping behavioural recovery for the patient (Pape, Senno, Guernon, & Kelly, 2005). The Sensory Modality Assessment and Rehabilitation Technique (SMART) was specifically designed for use in patients with UWS (Gill-Thwaites & Munday, 2011) and works by developing a rehabilitation programme for different levels of patient function. Finally, the CRS-R (Giacino et al., 2004) combines several clinical measures with a neuropsychological assessment of the patient. The measure is split into 23 items that measure a patient's responses to visual, auditory, motor, and communication tasks and their level of arousal. CRS-R has been

extensively validated in quantifying the level of consciousness in UWS and MCS patients (Giacino et al., 2004; Han et al., 2018; Sacco et al., 2011; Schnakers et al., 2008; Seel et al., 2010) and is the primary measure used in Chapter 6.

1.3.2.3. DOC Classifications

After performing a clinical assessment, patients are classified as one of several subtypes of DOC (Figure 1). Coma is the most severe type of DOC and is defined as having complete loss of stimulus-induced or spontaneous arousal. In cases of TBI, coma is typically the result of diffuse, bi-hemispheric lesions to the cortex, white matter, thalamus, or lesions to the paramedian tegmentum (Giacino et al., 2014). Patients in a coma lie with their eyes closed without any discernibly purposeful motor activity. They also show no sign of sleep-wake cycles on EEG. This state typically lasts around two-weeks before patients progress to UWS, MCS, or die (Giacino et al., 2014; Posner, Saper, Schiff, & Plum, 2007). Brain death is a related concept and has varying legal definitions in different countries (Haupt & Rudolf, 1999). In some countries, a patient is considered brain dead when the brain stem is completely deactivated (Bernat, 2006). Other countries require the entire brain (including the brain stem) to show no signs of activity (Bernat, 2006). Despite these differences, the clinical assessment of brain death is largely the same and tests for lack of any brainstem reflexes and continued demonstration of apnoea in persistently comatose patients (Laureys et al., 2004).

Diagnosis of UWS, formerly known as Vegetative State (Laureys et al., 2010), is made when spontaneous eye-opening remerges, a sign that the reticular activating system has at least partially recovered. Lesions in UWS patients are typically multifocal cortical contusions with diffuse axonal injury often accompanied with lesions to the thalamus (Giacino et al., 2014). The patient still does not have any discernible capacity to produce language or any other type of verbal or gestural communication. UWS patients also do not show any signs of reproducible purposeful behaviour in response to sensory stimuli (Bernat, 2006; Laureys et al., 2004). UWS is considered persistent (known as persistent vegetative state) when the patient has UWS for more than 1 month. After 3 months post non-TBI, persistent cases are often considered to be permanent. For TBI, patients are typically considered permanently in

UWS after 12 months. However, there are a small number of cases where patients wake up more than a year after injury (Choi et al., 1994; Dubroja, Valent, Miklic, & Kesak, 1995).

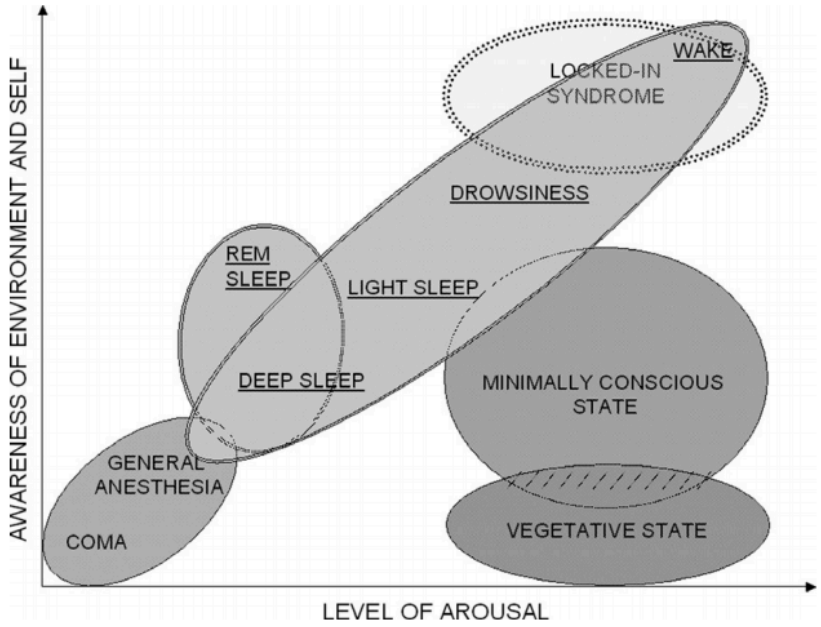

Figure 1. Correlation between arousal and awareness of the environment and self
This illustration shows the two main components of consciousness: arousal and awareness. It demonstrates various clinical and physiological states, like coma, anaesthesia, sleep and locked-in syndrome have correlated levels of arousal and awareness. States like UWS (shown here as Vegetative State) and MCS are dissociated from this correlation in that these patients have relatively high degrees of arousal but low to moderate levels of awareness. This figure is adapted from Laureys, Perrin, & Bredart (2007).

A patient is considered in a MCS when they recover to a chronic state of poor responsiveness to stimuli but shows some transient signs of awareness (Bernat, 2006; Giacino et al., 2002). The lesion profile of MCS patients is similar to that of UWS patients, but MCS patients often have fewer or no lesions to the thalamus. The key difference between MCS and UWS is that patients with MCS are capable of maintaining some measurable level of awareness, even if only for a short period. Recovery from MCS is much more common than recovery from UWS (Giacino et al., 2002). A patient is considered to be emerging from MCS (eMCS) when they regain the ability to communicate or functionally use objects (Laureys et al., 2004). A further sub-categorisation of MCS is the distinction between MCS+ and MCS- patients. MCS+ patients are capable of higher-level behavioural responses such as command following and some clear verbalisation. MCS- patients are capable of less complex neurocognitive functions like localising noxious stimulation and visual pursuit (Bruno, Vanhaudenhuyse, Thibaut, Moonen, & Laureys, 2011).

The term locked-in syndrome describes patients who are quadriplegic and anarthriaic due to damage to corticospinal and corticobulbar pathways but are nonetheless fully conscious (Laureys et al., 2004). Formally, it is defined as patients having sustained eye opening and awareness of their environment. They also have aphonia or hypophonia but are capable of vertical or lateral eye movement or blinking. They are capable of signalling yes or no with their eye movements, making it their primary form of communication (Laureys et al., 2004).

As previously discussed, providing accurate diagnoses and prognoses for DOC patients is a very challenging clinical problem, primarily due to the patient's inability to consistently respond to behavioural assessments. This has motivated research into identifying neuroimaging-based biomarkers that could be used to supplement behavioural assessments. This work has most recently focused on examining alterations in large-scale brain network connectivity in both DOC patients and healthy participants undergoing anaesthesia to characterise different levels of consciousness (Demertzi, Soddu, & Laureys, 2013; Di Perri et al., 2016; Fischer et al., 2016; MacDonald, Naci, MacDonald, & Owen, 2015). The following section outlines the imaging methodology and experimental literature surrounding large-scale brain network connectivity. The final section reviews the research applying brain network analyses to anaesthesia and patients with DOC.

1.3. Brain Networks

1.3.1. Introduction to Brain Connectivity

Modern theories in neuroscience emphasise the importance of dynamic interactions between distributed populations of neurons to enable flexible cognition and behaviour (Bressler & Menon, 2010; McIntosh, 1999; Mišić et al., 2016; Park & Friston, 2013). Advances in brain imaging methods have led to an explosion of research in this area, allowing for measurement of previously unobservable neurological processes. Concurrently, there have been significant advances in the statistical methods used to quantify changes in the brain during cognitive and behavioural tasks as well as during pharmacological and disease states. Specifically, modelling brain regions as nodes in a network where edges are represented by statistical relationships between these nodes

has resulted in a deeper understanding of brain structure and function (Mišić and Sporns, 2016).

The following sections detail the origins of brain connectivity and network analysis in neuroimaging, starting with a discussion of various imaging modalities. The neuroimaging modalities used in this thesis are all based on magnetic resonance imaging (MRI); therefore the following methodological description will discuss high-resolution T1-weighted MRI, functional MRI (fMRI), and diffusion MRI in detail. However, throughout the thesis, there will frequently be references to connectivity studies conducted using other imaging modalities, so they will also be briefly outlined. This will be followed by an overview of the initial functional connectivity studies conducted in humans, then a discussion on different large-scale brain networks and their role in cognitive processes.

1.3.2. Imaging Modalities for Brain Network Analysis

There are a variety of functional and structural imaging modalities that can be used to study brain network connectivity. These can be split into modalities that measure electrical currents and magnetic fields from the brain or that measure blood flow related to underlying neuronal activity.

EEG is a measurement of electrical potential differences between pairs of electrodes placed on the scalp. Initially discovered in 1929 by Berger, this technique is still widely used today due to its ability to provide high-temporal resolution brain activity data at a low cost (Bunge & Kahn, 2009). Specifically, EEG is a measure of transient electrical dipoles that are generated by the flow of electrical currents across neuronal cell membranes during depolarisation. EEG can be used as a measure of global brain dynamics that can measure electrical activity in different brain states, such as sleeping and waking. It can also be used to measure event-related potentials, a tool often used in cognitive neuroscience that averages the activity in response to specific triggers during behavioural paradigms, such as the presentation of a visual or auditory stimulus. EEG has been used extensively to study how anaesthesia affects brain activity (Akeju & Brown, 2017). EEG is also well suited for clinical use in patients with DOC due to its portability, cost-effectiveness, and lack of sensitivity to

metallic implants (Cruse et al., 2011). MEG has many similarities to EEG, but instead of measuring electrical activity directly, it measures the magnetic field produced by underlying electrical activity (Bunge & Kahn, 2009). Though it is less widely used than EEG, MEG has been previously used in patients with brain damage and DOC (Kotchoubey, Kaiser, Bostanov, & Lutzenberger, 2009).

Indirect measures of brain activity, like positron emission tomography (PET) and MRI, are also extensively used in both cognitive and clinical neuroscience. Though they are disadvantaged by the fact that they do not measure neuronal activity directly and have a comparatively low temporal resolution, they have great utility in their ability to measure brain activity with high spatial resolution. It is well established that regional brain activity elicits increases in regional cerebral blood flow (rCBF; Bunge & Kahn, 2009; Raichle, 1983). Modern PET imaging is based on pioneering work by Landau and colleagues, who demonstrated the first use of radioactive tracers to measure rCBF in animal models (Landau, Freygang, Roland, Sokoloff, & Kety, 1955). More precisely, this technique uses freely diffusible radiolabeled tracers injected into the bloodstream. The PET scanner detects pairs of gamma rays emitted by the tracer, ultimately producing 3-dimensional images of the tracer's concentration in different regions of the body (Bailey, Townsend, Valk, & Maisey, 2005). The most common tracer used in PET is Fludeoxyglucose F 18 (^{18}F-FDG). Uptake of ^{18}F-FDG is a measure of glucose uptake and is therefore highly correlated with tissue metabolism (Newberg, Alavi, & Reivich, 2002). A significant advantage of PET over similar imaging methods like fMRI is that using tracers like ^{18}F-FDG allow for the direct measurement of metabolic activity (Bunge & Kahn, 2009). It also allows for the synthesis of many different tracers that can mark a variety of targets in the brain like dopamine or serotonin receptors (C-11 or F-18 N-methylspiperone), opiate receptors (C-11 carfentanil), and microglial cell activation (translocator protein 18 kDa; Ito, Nyberg, Halldin, Lundkvist, & Farde, 1998; Tronel et al., 2017). PET imaging is widely used in the clinic and has been extensively used to aid in the diagnosis of patients with DOC (Di Perri et al., 2016; Stender et al., 2014).

The most widely used indirect measure of brain activity and the method used in this thesis is MRI. There are many types of MRI sequences that allow for

visualisation and quantification of different physiological processes (Glasser et al., 2016). This thesis focuses on high-resolution T1-weighted structural MRI, fMRI, and diffusion-weighted MRI (dwMRI). Each of these methods provides an important element in the study of cognitive and clinical neuroscience. Structural MRI allows for the examination of high-resolution neuroanatomy, fMRI detects changes in regional brain activity through associated changes in blood flow, and dwMRI examines the white matter connectivity between distributed brain regions (Bullmore & Sporns, 2009; Lerch et al., 2017).

MRI scanners use powerful magnetic fields, magnetic field gradients, and radio waves to produce images. The next section is a brief introduction to the physics behind MRI, followed by an overview of each sequence used in this thesis. Each atom contains protons that spin as a function of thermal energy. Proton spinning has angular momentum, due to an odd-numbered atomic mass. Because a proton is a charged particle, its spinning generates a looping electric current. This induces torque when placed in a magnetic field resulting in what is known as the magnetic moment (McRobbie & Moore, 2003). The body is largely made up of water molecules, each of which contains two hydrogen nuclei (that contain protons). When only weak magnetic fields are applied, these protons are randomly aligned, however when in a very strong magnetic field (like the one produced by an MRI scanner), the protons align to what is called the B_0 direction, which refers to direction and magnitude of the constant, homogenous magnetic field along the longitudinal axis. This results in each proton being aligned in one of two configurations, either parallel or anti-parallel to the B_0 direction. The majority of protons align to the B_0 direction, as this is the lowest energy state. When protons are in a magnetic field, they spin around their axis or as it is known in the relevant literature, are in precession. Precession can be thought of as a spinning top, where the proton wobbles but does not fall over its axis. The speed of precession is measured in how many times the proton spins about its axis per second and is known as the Lamor frequency (Currie, Hoggard, Craven, Hadjivassiliou, & Wilkinson, 2013). During scanning, a radio frequency pulse is applied, which excites the protons, causing them to flip from parallel to anti-parallel alignment. It is critical that this pulse has the same frequency as the Lamor frequency of the protons, a process known as resonance. The protons will then return back to their lowest energy

state. This is known as relaxation and causes changes in voltage of the receiver coils, resulting in a measurable signal. Using magnetic field gradients that linearly vary over space allows for different 2-D slices or 3-D volumes of the brain to be imaged. Different types of relaxation and contrasts result in different types of images produced by the scanner, which allows for the imaging of different anatomical and physiological processes (Currie et al., 2012).

T1-weighted structural images rely on T1 relaxation. A T1-weighted image is an image where the difference in signal intensity between tissues is dominated by differences in T1 relaxation time (Currie et al., 2012). T1-weighting is often used to acquire high-resolution structural images. In the context of brain connectivity analysis, these images are typically used for pre-processing and visualisation purposes. However, some recent work has shown that cortical thickness measured at different time points can act as a surrogate for connectivity in healthy and patient populations (Alexander-Bloch, Giedd, & Bullmore, 2013; Chen, He, Rosa-Neto, Germann, & Evans, 2008).

fMRI is dependent on the blood oxygenation level-dependent (BOLD) signal, an indirect measure of neuronal activity influenced by cerebral blood flow, cerebral blood volume, and the metabolic rate of oxygen consumption (Ogawa, Lee, Kay, & Tank, 1990; Ogawa, Menon, Kim, & Ugurbil, 1998). This measure relies on T2* relaxation, a process where protons fall out of phase in the x-y plane, resulting in a decrease in transverse magnetisation (Currie et al., 2012). Due to the BOLD response not being a direct measure of neuronal activity, many studies have sought to investigate the true neurobiological source of this signal (Logothetis, 2008). Previous work has shown the BOLD signal to be correlated with different measures of neuronal activity, including local field potentials (Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001) and gamma oscillations (Niessing et al., 2005). More recent work has combined fMRI with optogenetics to show that the BOLD signal correlates with optogenetically-driven neuronal activity (Kahn et al., 2013; Lee et al., 2010). Functional connectivity using fMRI is a powerful, commonly used method for examining the dynamics of large-scale brain networks. It involves correlating the BOLD signal from different brain regions across time. This can be done using *a priori*

defined regions of interest useful for specific hypothesis testing or at the whole-brain level (Bullmore & Sporns, 2009). Functional connectivity measures can then be compared between experimental conditions to identify differences in network dynamics as a function cognitive/behavioural tasks, disease state, or drug administration (Stamatakis, Orfanidou, & Papanicolaou, 2017).

In fMRI, and particularly in functional connectivity analyses, pre-processing methods are critical. There has been extensive work on understanding how to reduce variance due to non-neurobiological signals in fMRI data, with significant contributions in reducing motion artefacts (Power et al., 2018; Power, Barnes, Snyder, Schlaggar, & Petersen, 2012) and physiological noise (Birn, 2012; K. Murphy, Birn, & Bandettini, 2013). One area that has received less attention is the use of temporal filtering. Highpass filtering is critical, as it removes the low frequency drifts that are caused by the MRI scanner. However, lowpass filtering is more controversial as it has been shown to reduce task-related activity (Della-Maggiore, Chau, Peres-Neto, & McIntosh, 2002; Skudlarski, Constable, & Gore, 1999) and induce spurious autocorrelation during resting state analyses (Davey, Grayden, Egan, & Johnston, 2013). Several studies have sought to examine the best methods for temporal filtering directly. First, Murphy et al. (2009) showed that the use of global signal regression, a pre-processing step that averages the signal from the entire brain and regresses it from the data, significantly impacts functional connectivity results, even inducing anticorrelations between default mode network (DMN) and Task Positive Network (TPN) (Murphy et al., 2009). Further work by Chai et al. (2012) used an alternative pre-processing technique known as the CompCor method, in which only the signal from white matter and cerebrospinal fluid masks are regressed out, and showed that this technique reduces but not completely removes anticorrelations (Chai, Nieto Castañón, Ongür, & Whitfield-Gabrieli, 2012). This work suggests that anticorrelations may exist between brain networks, however, their physiological importance may have been exaggerated through global signal regression.

White matter tracts that are made up of bundles of neuronal axons directly connect different regions of the brain. Previously, these have been exclusively studied using histological techniques (Scannell, Burns, Hilgetag, O'Neil, & Young, 1999),

however using diffusion-weighted MRI allows for the measurement of white matter tracts *in vivo*. Diffusion Tensor Imaging (DTI) measures the restricted diffusion of water molecules in myelinated axons. Diffusion Spectrum Imaging (DSI) is similar to DTI, with the added benefit of measuring diffusion in multiple directions, allowing for the measurement of multiple groups of fibres in each location (Bullmore & Sporns, 2009). These images can then be used to measure direct structural connections between brain regions.

1.3.3. Large-Scale Brain Networks

Recent advances in cognitive and clinical neuroscience have suggested that interactions within and between distributed brain systems underlie complex and flexible cognitive processes (Bassett & Sporns, 2017). In this context, the term "large-scale" refers to these neural systems being distributed across the entire brain (Bressler & Menon, 2010). This is in contrast to smaller, microscale brain connectivity, which refers to the synaptic connectivity pattern of single neurons. This work has been conducted in the entire brain of *Caenorhabditis elegans* (Varshney, Chen, Paniagua, Hall, & Chklovskii, 2011; White & Southgate, 1986) and in regions of the mouse brain (Bock et al., 2011). However, assembling the entire human connectome at this scale still necessitates significant methodological advances and is therefore likely many years away (Sporns, Tononi, & Kotter, 2005).

As previously discussed, a variety of methods can be used to examine both structural and functional large-scale brain network connectivity. Over the past 25 years, the application of these methods has revealed a set of individual large-scale brain networks that play different roles in cognition and disease. Studies using various neuroimaging modalities and analysis methods, including both resting state and task-based analyses, suggest there are slight distinctions between the precise regions large-scale brain networks comprise of and their roles. However, the majority of studies find evidence of the existence of the following networks: The DMN, salience (SN) and attentional networks, Frontoparietal Control Network (FPCN), and sensorimotor networks (Bressler & Menon, 2010; Smith et al., 2009; Thomas Yeo et al., 2011).

Perhaps the most well known large-scale brain network, the emergence of the DMN began when researchers identified a set of brain regions that collectively decreased activity during goal-directed behavioural paradigms (Andreasen et al., 1995). These task-induced deactivations (TID) were typically thought to represent the cognitively unconstrained nature of resting baseline control conditions in neuroimaging experiments. However, meta-analyses of neuroimaging studies showed that activations were present across a wide range of experiments (Mazoyer et al., 2001; Shulman et al., 1997), suggesting that their presence was reflective of an underlying neurophysiological process and not a result of experimental noise.

In 2001, Gunsard, Raichle and colleagues published a series of studies focused directly on understanding resting state brain activity (Gusnard & Raichle, 2001). Resting state refers to a behavioural condition when the participant is inactive and quiet, yet fully conscious and aware. During resting state, participants are instructed to lay in the scanner with their eyes open, closed or fixated on a cross-hair and are asked not to think of anything in particular (Raichle, 2010). This condition was often used as a control condition in neuroimaging studies as a baseline of brain activity. In their review of the methodology, Gunsard and Raichle explain that concepts central to cognitive neuroscience, like activation and deactivation, and the control conditions used in experiments made interpretability and generalizability of neuroimaging results difficult. This motivated an in-depth study focused on defining a baseline of brain function (Gunsard & Raichle, 2001). Using PET to measure cerebral blood flow and oxygen consumption, the DMN displayed the highest level of activity during resting state. However, the oxygen extraction fraction, a ratio of blood flow and oxygen consumption, was hypothesised to remain consistent across the brain. This could then act as a baseline of brain activity where any deviation could be considered as "activity". Raichle et al. (2001) confirmed this hypothesis and showed that the regions making up the DMN, namely the PCC, medial prefrontal cortex, and bilateral angular gyrus, were not active as a result of a return to a control baseline, but represented an on-going physiological process in the brain. This series of studies defined this set of regions as the DMN (Figure 2A).

In parallel to the emergence of the DMN through PET imaging, fMRI had expanded to address questions regarding resting state connectivity. This began with demonstrations that the low frequency fMRI signal displayed synchronous activity in spatially distinct brain regions during resting state, forming networks (Biswal, Yetkin, Haughton, & Hyde, 1995; Friston, 1994). This method, termed functional connectivity (specifically seed-to-voxel functional connectivity), places seeds in regions-of-interest (ROI) and measures the average signal from those regions, then correlates that signal with all other brain areas. To directly examine the default mode hypothesis using fMRI, Greicius et al. (2003) defined seed regions in the posterior cingulate cortex (PCC) and ventral anterior cingulate cortex (vACC). The motivations for the selection of these seeds came from previous work showing these regions decreased activity during a cognitively demanding working memory task. They then measured functional connectivity of these seed regions during resting state. This analysis showed that the PCC, now known as a central node of the DMN, was strongly coupled with the vACC, inferior parietal cortex (IPC; including the angular gyri), medial orbitofrontal cortex, left dorsolateral prefrontal cortex (dlPFC), left inferolateral temporal cortex (ITC), and left parahippocampal gyrus (PHG). They also showed that this network displayed a similar spatial layout during a visual processing task, suggesting stability during passive sensory processing. Finally, the study showed inverse correlations between the PCC and three lateral prefrontal regions that were activated during the working memory task, suggesting the DMN was deactivated during cognitively demanding tasks (Greicius, Krasnow, Reiss, & Menon, 2003). Together, these findings demonstrated the existence of the DMN and laid the groundwork for future studies examining the interaction of the DMN and other large-scale brain networks.

Fox et al. (2005) directly examined the relationship between the DMN and what they defined as the TPN, which is now thought to comprise both Salience/attention networks[1] and FPCN (Figures 2B and C). Using resting state fMRI, they showed that regions of the DMN anticorrelated with regions of the TPN, suggesting that the two networks were dynamically opposed with one another, with

[1] The large-scale network literature has some inconsistency in naming conventions for attention networks and SN. For clarity,

the DMN increasing activity during internally directed attention and the TPN increasing activity during externally directed attention (Fox et al., 2005). This paper introduced a framework in which these two systems display competitive dynamics, leading subsequent work to adopt the idea that the DMN was not active during externally driven, attention-demanding tasks. This hypothesis has been expanded to suggest that the DMN is primarily involved in internally directed, introspective cognitive processes including long-term memory, prospective memory, theory of mind, and navigation (Buckner & Carroll, 2007). Spreng et al., (2010) used two goal-directed planning paradigms to show that the DMN was engaged with an autobiographical planning task, whereas the dorsal attention network (DAN) was involved in a visuospatial planning task. Furthermore, they demonstrated that both planning tasks engaged the FPCN and that the DMN coupled with FPCN during autobiographical planning. This study showed that the DMN is engaged during goal-directed cognition (Spreng, Stevens, Chamberlain, Gilmore, & Schacter, 2010).

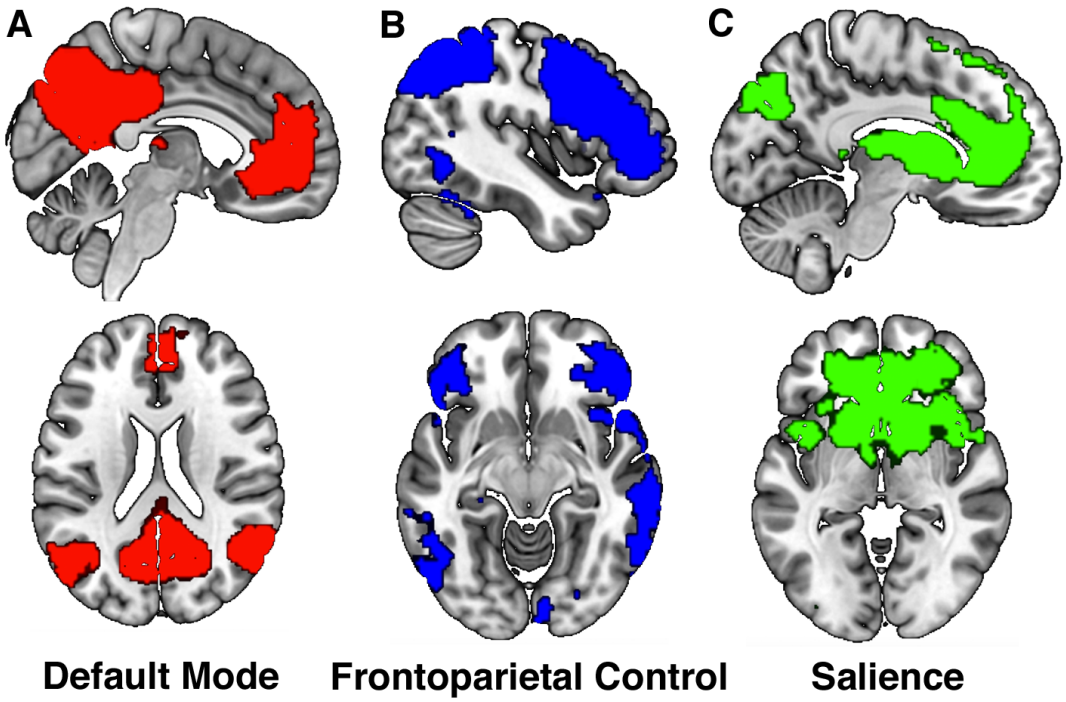

Figure 2. Examples of large-scale brain networks
A is the DMN, which includes PCC, medial prefrontal cortex and bilateral angular gyrus. **B** is the FPCN network, which includes dorsolateral prefrontal cortex, inferior frontal gyrus and supramarginal gyrus. **C** is the SN, which is comprised of the anterior cingulate cortex, anterior insula, lateral occipital cortex and intraparietal sulcus.

Recent work has also sought to measure DMN activity during externally oriented cognition. Vatansever et al. (2015) used graph theoretical measures to show that the DMN alters its connectivity pattern as task difficulty increases. Graph theory is a branch of mathematics that formally describes and analyses graphs. A graph is a set of nodes (vertices) and connections (edges) that provide an abstract representation of a real-world system or network (Bullmore and Sporns, 2009). In the context of brain networks, each node is a brain region, and each edge is a statistical relationship between those regions, often in the form of a correlation of BOLD signal or fibre count or density from a diffusion-weighted image. There are many graph theoretical measurements used to quantify various aspects of graphs (details of specific measurements used in this thesis are described in Chapter 4). Using an n-back working memory task, Vatansever et al. (2015) showed that DMN regions show altered topology as measured using graph theory. Specifically, they dynamically switch community membership and alter their nodal participation coefficient and strength when tasks become more difficult (Vatansever, Menon, Manktelow, Sahakian, & Stamatakis, 2015a). These findings suggest that the DMN does not deactivate during external tasks, but alters its pattern of connectivity in response to task demands. An additional study used multivoxel pattern analysis and a task-switching paradigm to show that regions within DMN encode task-relevant information during task performance (Crittenden, Mitchell, & Duncan, 2015). Another recent study used a cognitive flexibility task to show that the DMN plays a crucial role in autopilot behaviour, characterised by rapid selection of responses in predictable environments (Vatansever, Menon, & Stamatakis, 2017). An emerging theory is that the DMN is (or is part of) a global workspace, integrating information from across the brain to create a coherent subjective conscious experience (Dehaene, Changeux, & Naccache, 2011). This also fits with another theory of consciousness, integrated information theory (IIT). IIT presents a mathematical framework that describes a relationship between consciousness and its physical substrate, emphasising the importance of integrating information from disparate brain regions in a systematic way so as to make each subjective experience irreducible (Tononi, Boly, Massimini, & Koch, 2016). Primary nodes of the DMN are disproportionally large in humans (Leech & Sharp, 2014) and

are highly connected hub regions (Hagmann et al., 2008) making it a suitable candidate for system-wide information integration.

1.4. Neuroscience of Consciousness

Recent years have seen a significant increase in the quantity and quality of research in understanding the neural correlates of consciousness (NCC). The NCC are defined as the minimal neural mechanisms jointly sufficient for any one specific conscious experience (Crick & Koch, 1990; Rees, Kreiman, & Koch, 2002). There are two interpretations of the NCC depending on whether one is referring to the contents of consciousness or the overall state of conscious awareness, content-specific NCC and the full NCC (Koch et al., 2016). Content-specific NCC refers to the neuronal mechanisms associated with a specific conscious experience. These are the specific patterns of neuronal firing occurring during a subjective conscious experience. As an example, the content-specific NCC for visualising faces are the specific neuronal firing pattern that are activated whenever someone sees or imagines a face (Koch et al., 2016). In contrast to content-specific NCC, the full NCC are the complete neural substrates that support conscious experience in its entirety, irrespective of their contents. Put another way, they are the union of the sets of content-specific NCC for all possible experiential contents (Koch et al., 2016). Furthermore, a distinction can also be drawn between the full NCC and other factors that enable consciousness indirectly, like sufficient oxygen and glucose levels. This characterisation of consciousness is similar to differentiating between local and global states, where local states of consciousness are individuated perceptual experiences like imagery, bodily sensations, and affective experiences, whereas global states of consciousness refer to the overall conscious condition of an organism (Bayne, Hohwy, & Owen, 2016).

Many neuropsychological, clinical and pharmaceutical methods have been used to manipulate conscious experience with the aim of understanding its basis. As previously discussed, consciousness can be operationally defined as consistent maintenance of wakefulness and awareness (Laureys et al., 2004). This definition is somewhat distinct from Koch et al. (2016) in that it is based on specific behavioural markers so as to have clinical utility, and less concerned with identifying complete

neural patterns that correlate with subjective conscious experiences. In the healthy adult brain, the NCC are often studied by altering sensory awareness using neuropsychological tests while maintaining the level of wakefulness (Boly et al., 2013). These experiments use stimulus-manipulating cognitive and behavioural tasks like masking, binocular rivalry, and attentional blink. This work has shown that many cognitive processes occur without awareness and suggests some dissociation of consciousness from many higher-level cognitive functions (Boly et al., 2013; Dehaene, Charles, King, & Marti, 2014; van Gaal & Lamme, 2012).

While this work provides important evidence towards understanding cognition and some aspects of consciousness, fewer studies have focussed on directly studying different levels of wakefulness. The combination of anaesthesia, neuroimaging, and connectivity analyses provides the ability to systematically and reliably alter wakefulness while healthy participants perform tasks or undergo resting state scanning. Furthermore, studying patients with DOC using neuroimaging and connectivity analyses provides a unique opportunity to understand how brain damage can result in reduced levels of awareness and wakefulness. These pursuits not only allow for the study of consciousness but also help improve clinical care by deepening our understanding of neurobiology. The following sections will provide an overview of neuroimaging and brain connectivity studies, first with anaesthesia, and then with DOC, as models for examining the neurobiology underlying conscious experience.

1.4.1. Neuroimaging during Anaesthesia Administration

1.4.1.1 Activation Studies of Anaesthesia Administration

The earliest studies examining how anaesthetic drugs affect the brain focused on cortical activation, not connectivity (MacDonald et al., 2015). One study used isoflurane, an inhalational anaesthetic, combined with fMRI to show that both mild and moderate levels of sedation reduced BOLD response to tactile stimulation in the primary and secondary sensory cortices (Antognini, Buonocore, Disbrow, & Carstens, 1997). Another study used PET imaging to show that cerebral blood flow response to sensory stimulation was abolished at anaesthetic levels of propofol administration (Bonhomme et al., 2001). These studies appeared to suggest that any cortical reactivity

to sensory stimulation reduces as a function of the depth of unconsciousness. However, a study by Hofbauer et al. (2004) used noxious thermal stimulation at deep levels of propofol-induced sedation and showed preservation of cortical reactivity using PET (Hofbauer et al., 2004). This study demonstrated various brain functions are differentially affected by anaesthesia-induced unconsciousness (Heinke & Koelsch, 2005).

Other work in this area has expanded on the types of sensory stimulation studied while participants underwent sedation. For example, several studies have focused on how the brain responds to auditory stimulation and language processing while sedated (MacDonald et al., 2015). One study used low doses of sevoflurane and measured the activation to word stimuli relative to silence. Using fMRI, Kerssens et al. (2005) observed responses in bilateral superior temporal gyrus, bilateral parietal cortices, left frontal cortex, right occipital cortex, and right thalamus. However, when the sevoflurane dose was increased, activations in response to auditory stimuli were abolished (Kerssens et al., 2005). This suggests that auditory stimulation, like non-nociceptive sensory stimulation, is largely preserved at low levels of sedation, but is inhibited completely at deeper levels of unconsciousness.

An interesting aspect of examining the brain's response to auditory stimuli during sedation and anaesthesia is that this type of sensory stimulus can be made more cognitively complex by using various language paradigms. Davis et al. (2007) showed that bilateral temporal lobe responses to auditory, language-based stimulation are preserved in light propofol-induced anaesthesia. However, when comparing ambiguous versus non-ambiguous sentences, they found that higher-level, comprehension related activity associated with ambiguous sentences in inferior frontal cortex and posterior temporal regions was abolished (Davis et al., 2007). Another study showed that increased levels of sedation caused increased response time and error rates during a task where healthy participants needed to make perceptual or semantic decisions in response to auditory stimuli. They also found that sedation differentially affected how left inferior frontal gyrus (LIFG) and left inferior temporal gyrus (LITG) responded to semantic decision making (Adapa, Davis, Stamatakis, Absalom, & Menon, 2013). Specifically, they demonstrated that LIFG

activity, a region known to be involved in semantic processing, was reduced during sedation despite participants performing well on the semantic decision task. In contrast, activity in LITG was preserved and may therefore play a role in persistent semantic processing during sedation.

Taken together, these studies suggest that sensory information processing is impaired from more to less complex in a dose-dependent manner by anaesthesia (MacDonald et al., 2015). Activity is first impaired in higher-level prefrontal and association cortices at low doses, then reduces in lower-level sensorimotor regions as the dose increases. Activity in response to some types of sensory stimulation, particularly in response to nociceptive stimuli, appears to be preserved even at very deep levels of sedation and unconsciousness. These studies provide a crucial basis for our understanding of how anaesthetic drugs affect the brain. However, recent advances in both theory and methods in network neuroscience have revealed cognition to be a result of dynamic processing of a network of brain regions (Bullmore & Sporns, 2009; Mišić & Sporns, 2017). Therefore, activation studies identifying isolated brain areas can only provide a portion of the evidence necessary to understand how anaesthesia affects the brain. With this view in mind, the following section reviews the literature in applying network approaches to studying anaesthesia-induced unconsciousness.

1.4.1.2. Brain Connectivity During Anaesthesia-Induced Sedation

There is considerable evidence, mainly from resting state studies, that anaesthetic drugs alter functional connectivity both within and between well-known large-scale brain networks like DMN and FPCN (Hudetz, 2012; MacDonald et al., 2015; Mashour & Hudetz, 2018). Similar to the previously described activation studies, lower-level sensorimotor network intraconnectivity is largely preserved during sedation (Boveroux, Vanhaudenhuyse, & Phillips, 2010; Greicius et al., 2008; Martuzzi, Ramachandran, Qiu, Rajeevan, & Constable, 2011). Martuzzi et al. (2011) used sevoflurane and resting state functional connectivity to examine large-scale networks during sedation. They found that networks within auditory, sensorimotor, and visual cortices did not alter their connectivity patterns during sedation.

Connectivity within higher-level large-scale brain networks such as the DMN and SNs as well as interactions between these networks and sensorimotor networks is reduced at deeper levels of sedation (Boveroux et al., 2010; Greicius et al., 2008; Stamatakis, Adapa, Absalom, & Menon, 2010). One study showed that DMN connectivity persists during rest at very light levels of sedation (Greicius et al., 2008), however, as sedation increases, connectivity within and between DMN and executive control networks decreases. Furthermore, connectivity between these networks and the thalamus also decreases as sedation increases (Boveroux, Vanhaudenhuyse, & Phillips, 2010). Additional studies using task-based fMRI instead of resting state have also shown that functional connectivity is disrupted between primary auditory cortex and inferior frontal cortex in response to language-based auditory stimuli (Adapa et al., 2013; Liu et al., 2013). Stamatakis et al. (2010) used seed-based functional connectivity using the PCC as a seed region to show that at deeper levels of sedation, the PCC connects with regions not typically considered to be part of the DMN, namely the anterior cingulate (ACC), left premotor gyrus, and the pontine tegmentum. Connectivity within and between large-scale brain networks will be investigated further in Chapter 3.

Graph theoretical measures have also been applied to functional brain network dynamics during anaesthesia. Schrouff et al. (2011) used an information-theoretic method to quantify the loss of information integration during propofol-induced sedation. They found that in deep sedation, information integration was decreased within several large-scale brain networks, with areas becoming more temporally independent and therefore exchanging less information. This effect was most pronounced between frontal, occipital and parietal regions (Schrouff et al., 2011). Another study using dexmedetomidine with fMRI and whole brain graph theoretical measures showed that both local and global efficiency were reduced during sedation. They also showed that dexmedetomidine reduced overall strength of connectivity without altering degree distribution (Hashmi et al., 2017). The application of network science to studying brain connectivity during anaesthesia is interesting, however, its utility in predicting conscious state has not been fully assessed. This will be investigated more extensively in Chapter 4.

Another interesting use of sedation is to model structure-function relationships in brain networks. This is a particularly interesting and important question, as it is unknown how a dynamic, complex pattern of functional connectivity can arise from a largely static underlying structural connectivity (Honey et al., 2009; Park & Friston, 2013). Barttfeld et al. (2015) used fMRI and a non-human primate model to understand how the relationship between brain structure and function changes during anaesthesia. They found that during deep levels of sedation, there was significantly greater overlap between functional and structural brain networks suggesting that as sedation increases functional connectivity reduces in complexity and inherits the pattern of its underlying structural connections. Furthermore, Barttfeld et al. (2015) found that functional networks have reduced small-world properties and far fewer anticorrelations. This reflects the finding by Stamatakis et al. (2010) that during sedation there is an increase in connectivity between regions not normally connected during wakefulness. The relationship between brain network structure and function has not been directly examined in humans before. Chapter 4 provides initial evidence that similar structure-function patterns can also be observed in humans.

1.4.2. Neuroimaging in Disorders of Consciousness

Behavioural assessment of patients with DOC is a critical element of the diagnostic and prognostic procedure. While it remains the clinical gold standard, neuroimaging approaches represent a significant opportunity to improve the diagnostic process by identifying biomarkers specific to various conscious states. The most sophisticated approaches combine neuroimaging with modern machine learning techniques with the aim of generalising and automating the evaluative process. Furthermore, well-designed neuroimaging experiments in patients with DOC expand our knowledge of the underlying NCC. The following section outlines the literature regarding neuroimaging in DOC patients and highlights areas where more work is needed.

1.4.2.1. Structural Brain Imaging in DOC

In acute TBI, computerised tomography (CT) scans are typically used to assess initial damage (Giacino et al., 2014). This is due to the accessibility and speed of acquiring a

CT scan as opposed to an MRI scan and because CT scans are more sensitive to haemorrhage or lesions that require immediate surgery (Smith, Citerio, & Kofke, 2016). After patients stabilise, MRI is the best method to provide high-resolution images of the extent and location of damage. Though standard T1-weighted images provide some information for clinicians, they do not allow differentiation between the various DOC diagnoses. However, some work has shown that the preservation of certain brain regions visible on T1-weighted images are predictive of recovery from UWS. Kampfl et al. (1998) followed a group of 80 patients with UWS[2] from 2 to 12 months after injury. High-resolution T1-weighted images were assessed separately by three neuroradiologists who recorded the number, signal intensity, size, and location of lesions. They found that patients with persistent UWS (i.e. those who did not recover) were significantly more likely to have lesions in the corpus callosum, dorsolateral brainstem, and corona radiata than patients who recovered, suggesting a critical role for brainstem and white matter integrity in recovery (Kampfl et al., 1998). As such, DTI methods have also been applied to directly assess white matter integrity following DOC. The application of these methods has proven very fruitful, having shown to perform better than various clinical markers in recovery prediction in cohorts of patients with DOC from both anoxic (Luyt et al., 2012) and TBI (Galanaud et al., 2012). DTI has also been used to assess the integrity of specific large-scale brain networks as markers of DOC diagnosis and recovery. Fernández-Espejo et al. (2011) assessed the structural connectivity of the DMN in 52 patients with UWS, MCS, or eMCS. They found that DMN structural connectivity, as well as structural connectivity between PCC and the thalamus, was less preserved in patients with lower levels of consciousness. They also found that integrity of connections in posterior regions of the DMN correlated significantly with patients' behavioural signs of awareness with integrity lowest in UWS patients. These findings suggest that specific DMN connections are essential in assessing DOC patients and demonstrate the utility of focusing on large-scale brain network connectivity to aid in diagnosis (Fernández-Espejo et al., 2011).

[2] UWS is referred to as VS in the paper.

1.4.2.2. Functional Brain Imaging in DOC

Many functional brain imaging modalities including fMRI, PET, and electrophysiological techniques have been used to assess patients with DOC. The goals of these studies are twofold. First, they combine neuroimaging with task paradigms to identify preserved cortical processing or volitional capacity in the form of command following. Second, they aim to develop a tool sensitive to a biomarker for a specific diagnosis of an individual patient.

An early case report from de Jong et al. (1997) examined changes in rCBF in a TBI patient in persistent UWS. They presented auditory stimuli in the form of a story told by the patient's mother and found activation in rostral anterior cingulate, premotor cortices, and right middle temporal gyrus (de Jong, Willemsen, & Paans, 1997). A critical study by Menon et al. (1998) examined a patient in persistent UWS using PET imaging. They used a face-recognition paradigm with familiar faces to the patient as well as scrambled versions of these images as a control. They found activation in response to faces in the right fusiform gyrus that spread to the dorsal cerebellum and extrastriate areas 18 and 19. Two months after the scanning was performed and six months after illness onset, the patient became more responsive and came to clearly recognise faces and produce short sentences (Menon et al., 1998). Other studies have also examined activation of brain regions in response to other types of sensory stimuli, including auditory (Boly et al., 2004; Laureys et al., 2000) and noxious somatosensory stimuli (Laureys et al., 2002).

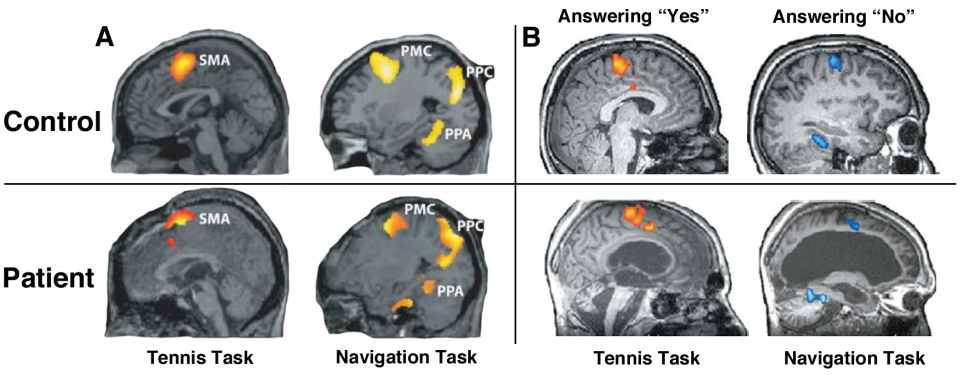

Figure 3. Imagining playing tennis to communicate responses
A shows significant activations for healthy control (top) and a patient with UWS in response to mental imagery paradigms (Tennis task on the left and Navigation task on the right; Adapted from Owen et al., 2006). **B** shows a healthy control and a DOC patient using the Tennis task and Navigation task to respond yes or no to simple questions (adapted from Monti et al., 2010).

This work motivated further research into "active" paradigms aimed to identify motor-independent signs of command following in DOC patients (Giacino et al., 2014). These studies use a variety of cognitive tasks using visual, visuospatial, and motor imagery tasks to identify covert awareness (Boly et al., 2007). In a seminal study by Owen et al. (2006), a patient suffering from UWS was asked to imagine playing tennis or moving around their home (i.e. a spatial navigation task) while undergoing fMRI scanning. Owen et al. (2006) observed activation in the supplementary motor area (SMA) when the patient imagined playing tennis and activation in the posterior parietal cortex (PPC), premotor cortex (PMC), and parahippocampal gyrus (PPA) during the imagined spatial navigation task. Both of these task responses were very similar to healthy control participants, suggesting the patient in UWS was capable of modulating their brain activity in response to task instructions (Owen et al., 2006; Figure 3A). A subsequent study expanded these tasks to a group of 54 patients with DOC. Five patients showed activity that suggested they were able to modulate their brain activity in response to these behavioural paradigms. This study took these results a step further by using the response to the tasks to ask yes or no questions to the patients, like whether or not they had a brother or sister (Figure 3B). Of the five responding patients, one was able to answer these yes or no questions. Importantly, this patient was incapable of communicating in any way at the bedside (Monti et al., 2010). An additional study by Chennu et al. (2014) used EEG and graph theoretical measures to develop a novel topographical metric known as modular span. Graph theoretical metrics like path length, modularity, and clustering all quantify topological aspects of networks but do not capture their topographical characteristics. Modular span measures the average weighted topographical distance over the scalp, spanned by a specific network module. This novel metric allowed Chennu et al. (2014) to assess long-range connectivity in different frequency bands directly. This measure can then be visualised to demonstrate differences between healthy controls and patients (Figure 4). Their results demonstrated that patients who were able to respond to a volitional paradigm during scanning had more alpha network long-range connectivity (similar to healthy controls) than DOC patients who were unable to respond.

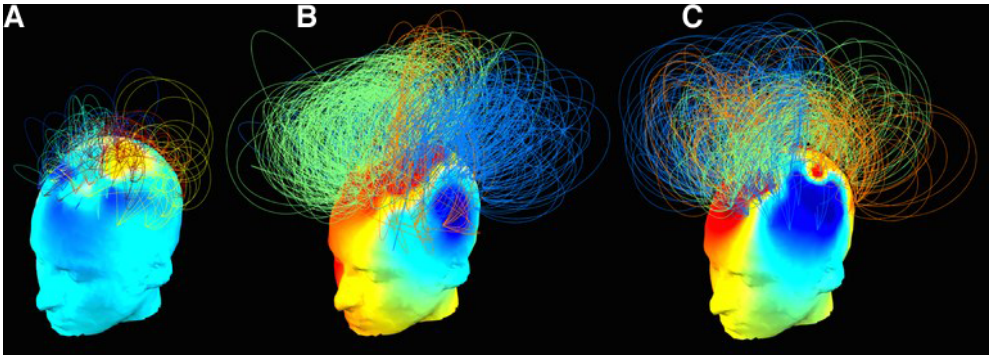

Figure 4. Spectral signatures measure with EEG during Tennis Task
Panel **A** shows a UWS patient incapable of responding to the task, where as panel **B** shows a UWS patient responding to the task. Panel **C** shows a pattern in a healthy control. Adapted from Chennu et al. (2014).

This work has since been expanded to additional tasks and analysis techniques to identify wilful modulation of brain activity in DOC patients (Bardin, Schiff, & Voss, 2012; Bekinschtein, Manes, Villarreal, Owen, & Della-Maggiore, 2011; Cruse et al., 2011; Monti, Pickard, & Owen, 2012). Cruse et al. (2011) extended this concept to incorporate the use of EEG instead of fMRI. EEG has the advantage of being significantly cheaper and portable, making it possible to assess awareness at the patient's bedside. They asked patients to imagine moving their right hand and toes in response to a command and found that three of the 16 patients studied could reliably generate significantly different EEG responses to these commands. The EEG responses these patients produced were similar to those produced by control subjects in response to the same commands (Cruse et al., 2011). Naci et al. (2014) used healthy participants and one DOC patient with a novel movie-watching task found a common neural signal underlying similar experiences that could then be used to detect these experiences in the absence of behaviour. The movie used was a scene from an Alfred Hitchcock film in which a young boy finds a revolver and pulls the trigger repeatedly without realising it is loaded. During viewing, healthy subjects showed synchronised activity in auditory, visual, and frontoparietal regions known to support executive function and cognitive control. This method was then applied to DOC patients, and the authors found that one patient who had remained unresponsive for 16 years displayed highly synchronised activity in the same regions as the healthy subjects in response to the movie (Naci, Cusack, Anello, & Owen, 2014). Taken together, these studies demonstrate that behavioural assessment of DOC patients is not sufficient to establish whether they have preserved mental capacity despite being in an altered state of consciousness. It is important to note, however, that not all patients who show

behavioural signs of awareness show positive fMRI results and that patients with aphasia and other cognitive deficits, but preserved awareness, may not be detected with these methods (Majerus, Bruno, Schnakers, & Giacino, 2009).

In addition to the studies examining volitional capacity, other neuroimaging work has aimed to identify biomarkers associated with different sub-categories of DOC (Aubinet et al., 2018; Giacino et al., 2014). This work typically uses resting state fMRI or FDG-PET, which, as previously discussed, is thought to be a measure of spontaneous neuronal activity in the absence of any behavioural or cognitive paradigm. As such, resting state scanning is particularly apt for non-responsive patients with DOC. The use of resting state allows for the examination of alterations in large-scale network connectivity and how these can be used as biomarkers for DOC diagnosis (Demertzi et al., 2013). PET imaging shows an overall reduction in brain metabolism in DOC patients typically between 40-50% (Laureys & Schiff, 2012). There is also evidence of reduced interhemispheric connectivity during resting state scanning in DOC patients (Ovadia-Caro et al., 2012). Early studies using PET imaging with voxel-based analyses show reduced activity in medial frontoparietal regions and suggest that recovery of activity in the DMN correlates with recovery of awareness (Laureys et al., 1999). As previously discussed, the DMN is characterised by its role in mediating attention to internal states while the FPCN is important in processing externally directed attention and cognitive control (Vanhaudenhuyse et al., 2010). More recent work has further demonstrated the importance of the DMN in DOC patients' recovery by showing that DMN connectivity is absent in patients considered brain dead but is mostly preserved in UWS and MCS patients (Boly et al., 2009; Cauda et al., 2009; Vanhaudenhuyse et al., 2010).

A recent study by Di Perri et al. (2016) took network analyses a step further by examining how the DMN and FPCN network interact with each other in DOC patients. They used resting state fMRI and PET imaging in a large cohort of DOC patients ranging from UWS to eMCS to determine whether within-network DMN and anti-DMN (i.e. task positive or FPCN; networks typically anticorrelated with DMN) connectivity can predict the degree of disease severity. Due to previous work in healthy participants that found persistent anticorrelations between these networks at

rest (Fox et al., 2005), they hypothesised that patients who were less conscious would have reduced anticorrelations. Indeed, the authors found that UWS and MCS patients had reduced within-DMN connectivity (Figure 5A) and positive correlation between these two brain networks while eMCS patients had, like the healthy controls, anticorrelations between these networks (Figure 5B). There was also a significant difference in the degree of anticorrelation between the eMCS patients and the healthy controls. Di Perri et al. (2016) also showed that this increase in anticorrelations was mirrored by increased brain metabolism measured using ^{18}F-FDG PET. This increase in metabolism potentially underlies the increased taxing of neuronal systems to maintain the necessary balance between information integration and segregation (Baars, 2002; Deco, Tononi, Boly, & Kringelbach, 2015). Together, these results suggest that this dynamic switching between modes of activity from internally-directed to externally-directed attention are essential for normal cognitive processing (Chennu, Stamatakis, & Menon, 2016). This finding highlights the importance of analysing both intra- and inter-network connectivity in DOC patients and provides a potential neuroimaging-based biomarker for differentiating between MCS and eMCS patients.

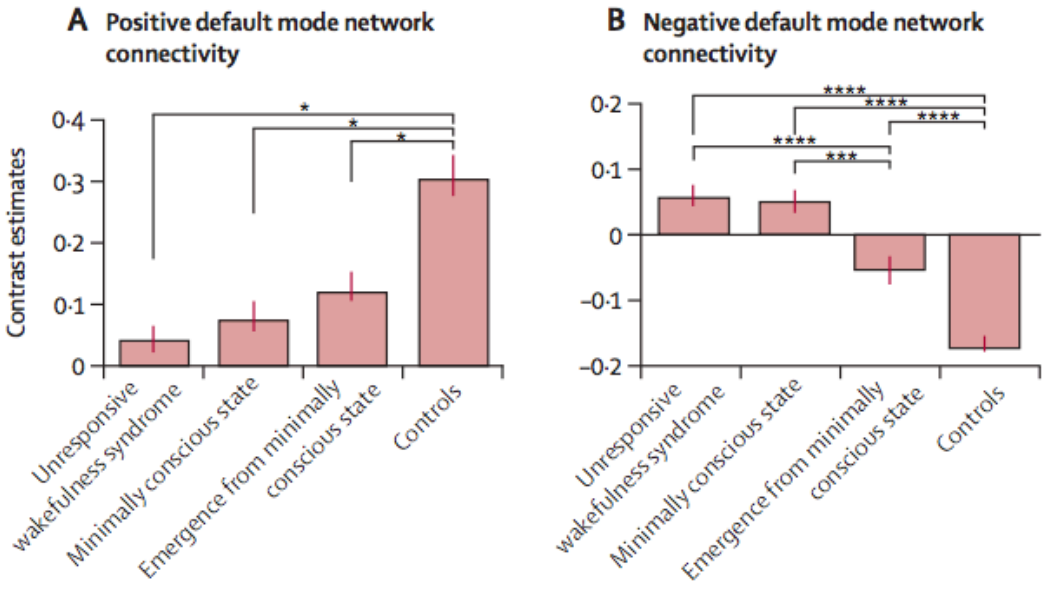

Figure 5. DMN connectivity and brain metabolism in DOC
A shows that within DMN connectivity increases from lower to higher levels of consciousness, **B** demonstrates that anticorrelations with DMN (termed negative DMN connectivity in this study) decrease at lower levels of consciousness. Adapted from Di Perri et al. (2016).

In addition to examining between-network interactions, the relationship between brain network structure and function has also recently been examined. Demertzi et al., (2019) demonstrated that patients with DOC had less complex functional brain network connectivity than healthy controls. Furthermore, they found that this less complex pattern overlapped significantly with structural connectivity defined using DWI. These results suggest that, like propofol-induced anaesthesia (Barttfeld et al., 2015), DOC patients have less complex patterns of functional connectivity that are largely restricted to anatomically connected regions.

Given that these analyses are all performed at the group level, and that misdiagnosis by behavioural assessment is so high, there has been some effort in developing machine learning methods to improve diagnosis and provide clinical information regarding DOC patients (Noirhomme, Brecheisen, Lesenfants, Antonopoulos, & Laureys, 2017). The applications of machine learning to neuroimaging in DOC patients, like neuroimaging in DOC in general, fall into two separate groups. One group of applications focus on classification of active behavioural paradigms and another group of applications provides a diagnostic estimate. Bardin et al., (2012) used multivariate pattern classification to distinguish between patients who were able to modulate their brain activity in response to a command following paradigm. This analysis was performed within each subject, and the authors found that the responses of 3 out of 5 patients could be accurately classified. Furthermore, they showed that in two of these patients, responses to a yes or no question could be differentiated (Bardin et al., 2012). This study demonstrated that multivariate modelling of brain data combined with a machine learning classifier is capable of detecting changes in brain activity in DOC patients and may provide more sensitivity than traditional activation approaches.

There have also been a considerable number of studies involving the development of classifiers that aim to provide a diagnostic estimate (Chennu et al., 2017; Demertzi et al., 2015; Fernández-espejo et al., 2011; King et al., 2013; Phillips et al., 2011). Demertzi et al. (2015) used resting state fMRI to measure functional connectivity within several large-scale brain networks in 73 patients with MCS, UWS, and coma. They found that using network features as input to support a vector

machine classifier resulted in high classification accuracy for each large-scale brain network with auditory network providing the best discrimination (Demertzi et al., 2015). Another study used EEG as an input feature to develop a machine learning model for diagnosing DOC patients. Chennu et al. (2017) used graph theoretical measures to model spectral connectivity from high-density EEG measured at rest in 104 DOC patients. They found that their classifier could predict clinical diagnosis, brain metabolism, and clinical outcome after one year in individual patients. Furthermore, they found that several patients who were initially misdiagnosed by behavioural analysis were correctly diagnosed by their algorithm (Chennu et al., 2017). These studies show the importance and utility of further developing machine learning methods to aid in the diagnosis and prognosis of DOC patients. The rapid development of machine learning more generally allows for cutting-edge algorithms to be tested in the context of neuroimaging-based assessment of DOC and is the focus of Chapters 5 and 6.

1.5. Thesis Outline

The previous sections have outlined general anaesthesia, DOC and the current state of knowledge of the neuroscience of consciousness. Despite the large and growing body of work in this field, many questions remain to be answered, ranging from improvement in basic methodology and a deeper characterisation of network dynamics during drug-induced and pathological states of unconsciousness to the application of advanced statistical methods and machine learning to improve DOC patient diagnosis and prognosis. Having identified these evidential gaps in the current literature, this thesis reports on several experiments aimed at understanding different aspects of consciousness in the context of clinical neuroscience. Chapter 2 examines a specific aspect of pre-processing of fMRI data to improve measurements of intra- and inter-network connectivity that is then applied to the analyses reported in the remaining chapters. Chapter 3 uses a multivariate statistical analysis of brain network connectivity to understand how brain networks interact during propofol sedation. Chapter 4 examines the relationship between these brain network interactions and their underlying structural connections. Chapter 5 develops a brain network-based machine learning algorithm to differentiate between the performance of a mental

imagery task and rest and validates it in healthy subjects. Chapter 6 applies this network-based machine learning algorithm to patients with DOC to understand how large-scale brain network connectivity can be used to aid in clinical assessment. Chapter 7 discusses the conclusions that can be drawn from all experiments, highlights how the results expand the current state of knowledge about the neuroscience of consciousness and suggests directions for future research.

Chapter 2. Spectral diversity in DMN connectivity reflects behavioural state

2.1. Introduction

Large-scale brain networks form a complex system of connections capable of supporting a wide range of behaviours. Using fMRI, multiple brain networks have been consistently identified using a variety of statistical techniques (Smith et al., 2009; Thomas Yeo et al., 2011). As previously discussed, the DMN, despite initial studies suggesting it deactivates during externally driven tasks (Mazoyer et al., 2001; Shulman et al., 1997), has recently been shown to play a crucial role in a variety of introspective and attentionally demanding tasks (Buckner & Carroll, 2007; Vatansever et al., 2015a, 2017). This includes interoceptive attentional processes such as autobiographical memory, prospective memory, moral judgement, and theory of mind (Andrews-Hanna, Reidler, Sepulcre, Poulin, & Buckner, 2010; Buckner, Andrews-Hanna, & Schacter, 2008; Spreng, Stevens, Chamberlain, Gilmore, & Schacter, 2010) as well as attentionally demanding executive function tasks (Vatansever et al., 2015a, 2017).

A recurring element in brain network studies is that network fluctuation is maximally observed at low frequencies. Therefore, most researchers bandpass filter their data prior to statistical analysis, typically retaining frequencies above ~0.01Hz and below ~0.1Hz (Biswal et al., 1995; Fox et al., 2005; Fransson, 2005; Greicius, Krasnow, Boyett-Anderson, et al., 2003; Murphy et al., 2013). While it is true that most statistical power is contained within these low frequencies, recent work has shown that higher frequencies (those above 0.1Hz) also contain non-noise signal (Kalcher et al., 2014). Highpass filtering is relatively uncontroversial, as it removes low frequency drifts due to scanner noise. However, lowpass filtering has been shown to be problematic in that it decreases sensitivity to task-related activations without increasing specificity (Della-Maggiore et al., 2002; Skudlarski et al., 1999) and induces spurious autocorrelation in resting state studies (Davey et al., 2013). Furthermore, recent work has found that many artefactual signals have spectral peaks within the low frequencies used to identify large-scale networks or are aliased into these frequencies

by the low sampling rate of fMRI (Birn, Smith, Jones, & Bandettini, 2008; Van Dijk et al., 2018). Recently, these problems have been dealt with through novel denoising statistical techniques such as anatomical CompCor (Behzadi, Restom, Liau, & Liu, 2007; Chai et al., 2012), in which principal components from white matter and CSF and movement parameters and their first derivatives are removed.

There is a growing body of literature focused on understanding high frequency contributions to functional connectivity measures (Boyacioglu, Beckmann, & Barth, 2013; Chen & Glover, 2015; De Domenico, Sasai, & Arenas, 2016; Gohel & Biswal, 2015; Kalcher et al., 2014; Lewis, Setsompop, Rosen, & Polimeni, 2016; Wu et al., 2008). A recent study by Chen and Glover (2015) collected resting state data at different echo times (TE) to examine the relative contributions of BOLD and non-BOLD components to resting state connectivity at different time scales. Their study focused on connectivity in DMN and the executive control network (a subnetwork of the DAN) and found altered spatial patterns at different frequency bands. Notably, they found that some subjects showed anticorrelations between DMN and executive control network only in low frequencies. These connections became positive when sampled between 0.2-0.4 Hz, suggesting that anticorrelations between DMN and DAN, as initially described by Fox et al., (2005) and Fransson (2005), may be frequency dependent. Furthermore, studies comparing network connectivity between resting state and tasks tend to retain high frequencies with the assumption that behaviourally relevant signals are in the higher ranges (Cole, Bassett, Power, Braver, & Petersen, 2014; Cole, Ito, Bassett, & Schultz, 2016; Schultz & Cole, 2016). This is because univariate analyses of task-based fMRI data typically employ highpass filters. Though this is a sensible assumption, a direct investigation of bandpass and highpass filtered datasets and the effect of filtering on DMN connectivity during tasks has not yet been performed.

The current study fills this gap in the literature by presenting novel evidence of higher frequency DMN connectivity and its relationship to behaviour. Data from resting state, a finger opposition task (blocked design), and a Stop Signal Task (SST; event-related design) were temporally preprocessed in two different ways, one using bandpass filtering ($0.009 < f < 0.08$) and one using highpass filtering ($0.009 < f < 0.25$).

These tasks were selected to investigate how temporal filtering affects DMN connectivity during both block and event-related task designs. We hypothesised that the inclusion of high frequencies would alter connectivity patterns between DMN and DAN across behavioural state. Furthermore, we hypothesised that retaining high frequency fluctuations in the DMN would be associated with task performance. This work was published in the *Journal of Cognitive Neuroscience (*doi: 10.1162/jocn_a_01213)

2.2. Methods

2.2.1. Participants

This study was approved by the Cambridgeshire 2 Research Ethics Committee (LREC 08/H0308/246) and all participants gave written informed consent before testing. All participants were right-handed, had no history of psychiatric or neurological disease, no history of drug or alcohol abuse, no contraindication to MRI scanning or severe claustrophobia, and were not taking medication that affected their physical or cognitive performance. Additional criteria required a score above 70 on the National Adult Reading Test (NART) and a score above 23 on the Mini Mental State Exam (MMSE). Twenty-two healthy participants were included in the study (19-57 years old, mean = 35.0, SD = 11.2, 9 females and 13 males). Participants had an average score of 117.1 (SD = 5.76) on the NART and 29.33 (SD = 0.85) on the MMSE.

2.2.2. Paradigm Specifications

2.2.2.1. Resting State

Participants underwent a 5-minute resting state scan during which they were instructed to stay still and keep their eyes closed.

2.2.2.2. Motor task

The motor task is a boxcar design, self-paced finger opposition paradigm with 5 alternating 30-second blocks of task and fixation (5 minutes in total). Participants were instructed to touch their fingers with their right thumb, moving sequentially

from the index to the little finger. They continued this process throughout the entire duration of the task block. The task blocks were initiated by participants seeing the word "move" on the screen while fixation blocks started with participants seeing the word "rest".

2.2.2.3. Stop Signal Task

The SST is an event-related fMRI paradigm and has previously been described by Rubia et al. (2003; 2007). Briefly, an arrow facing either left or right is displayed on a computer screen (white on black background). Participants are instructed to respond with their right index (left-facing arrow) or middle (right-facing arrow) fingers depending on the arrow direction (go trials). There were a minimum of 3 and a maximum of 7 go trials between each stop trial. The stimulus duration for go trials was 1000ms while the stop trial stimulus duration was a minimum of 100ms and a maximum of 300ms. In a subset of trials, the arrow is followed by another arrow pointing upwards (stop trials), indicating the participant should stop themselves from pressing a button. For the first stop trial, the interval between onset of the go trial and onset of the stop trial (i.e. stop signal delay; SSD) was 150ms. This task has a total of 240 trials including 200 go trials and 40 stop trials and is designed to alter the time interval between go and stop signals so that each participant succeeds in 50% of the stop trials. Successful inhibition resulted in the SSD to increase by 50ms, whereas unsuccessful inhibition resulted in SSD decreasing by 50ms. This ensures that each participant is working at the edge of their ability, resulting in a comparable level of difficulty between subjects (Rubia, Smith, Brammer, & Taylor, 2003; Rubia, Smith, Taylor, & Brammer, 2007). Of the four conditions included in this task (go success, go failure, stop success, stop failure), only go success trials were considered in this study. This is because we are interested in measuring task-related activity generally, not neural activity specific to stop inhibition. One subject was removed from the stop signal task because of excessive head motion during scanning leaving 21 subjects in this group.

2.2.3. Data Acquisition

MRI data was obtained using a Siemens Trio 3T scanner (Erlangen, Germany) equipped with a 12-channel head matrix transmit-receive coil at the Wolfson Brain Imaging Centre at Addenbrooke's Hospital in Cambridge, UK. First, participants underwent a high resolution T1-weighted, magnetization-prepared 180 degree radio-frequency pulses and rapid gradient-echo (MPRAGE) structural scan (TR = 2300ms; TE = 2.98 ms; TA = 9.14 min; flip angle = 9°; field of view (FOV) read = 256mm; voxel size = 1.0 x 1.0 x 1.0mm, slices per slab = 176). The structural scan was followed by whole-brain echo planar imaging (EPI) for the resting-state, motor task, and SST (TR = 2000ms; TE = 30ms; flip angle = 78°; FOV read = 192 mm; voxel size = 3.0 x 3.0 x 3.75mm; volumes = 160; slices per volume = 32). Other tasks were also completed by the participants (N-back, Tower of London and Rapid Visual Information Processing) and were presented in a randomised order. The motor task always followed the resting state scan and randomisation of the remaining tasks was used to avoid potential fatigue on any single task. Portions of this data were used in previously published studies to address alternative experimental questions (Vatansever, Manktelow, Sahakian, Menon, & Stamatakis, 2016; Vatansever et al., 2015a; Vatansever, Menon, Manktelow, Sahakian, & Stamatakis, 2015b).

2.2.4. Pre-processing

Pre-processing for resting state functional data and task functional data followed an identical pipeline described below.

2.2.4.1. Spatial Pre-processing

Spatial pre-processing for functional and structural images was performed using Statistical Parametric Mapping (SPM) 8.0 (http://www.fil.ion.ucl.ac.uk/spm/) and MATLAB Version 2008a (http://www.mathworks.co.uk/products/matlab/). A previous study performed a comprehensive analysis of pre-processing step order and found the following order to be most optimal: 1) Slice timing correction, 2) Realignment, 3) Spatial normalization, 4) Smoothing, and 5) Noise Signal removal (Weissenbacher et al., 2009). We deviated slightly from step 5 in that we used the anatomical CompCor method for removing noise signals (Behzadi et al., 2007).

Otherwise, we used the pre-processing step order from Weissenbacher et al. (2009) defined in their paper. For all functional scans, the first 5 volumes were removed to eliminate saturation effects and achieve steady-state magnetization. Subsequently, functional data were slice-time adjusted and underwent motion correction using SPM realign. High-resolution T1 structural images were co-registered with the mean EPI and segmented into grey matter, white matter, and cerebrospinal fluid masks and were spatially normalised to Montreal Neurological Institute (MNI) space with a resolution of 2 x 2 x 2mm cubic voxels (Ashburner & Friston, 2005). Transformation parameters were applied to motion corrected functional images that were then smoothed with a 6 mm FWHM Gaussian kernel.

2.2.4.2. Temporal filtering

To deal with physiological and movement-related noise, data underwent despiking with a hyperbolic tangent squashing function followed by the anatomical CompCor (aCompCor) technique. aCompcor removes the first 5 principal components of the signal from white matter and cerebrospinal fluid masks as well as the motion parameters, their first-order temporal derivatives, and a linear detrending term (Behzadi et al., 2007). aCompcor has been shown to effectively remove physiological noise components that would otherwise be aliased into data sampled at the standard ~2s TR. This process was identical for both temporal filtering sets.

Each subject's preprocessed functional images were then highpass (0.009 Hz < f < 0.25 Hz) and bandpass filtered (0.009 Hz < f < 0.08 Hz). 0.25 Hz was selected as the low-pass filter for the high-pass group because it is the Nyquist frequency for data acquired with a TR of 2s. The Nyquist frequency is the highest reliably sampled frequency given the sampling rate of the data (i.e. half the sampling rate). This resulted in two groups of functional images that were identical except that one was highpass filtered and the other was bandpass filtered. Temporal filtering was performed using a Fast Fourier Transform (FFT).

2.2.5. Obtaining Power Spectrum

We calculated the power spectrum at the single subject level using the Fast Fourier Transform (FFT) implemented with MATLAB Version 2008a

(http://www.mathworks.co.uk/products/matlab/). The frequency bands were restricted to signals between 0.009 and 0.25 Hz.

2.2.6. ROI Definition

ROIs were defined as 10mm spheres and centred at seven regions (Fox et al., 2005). Four of the regions made up the DMN and included the PCC (-6, -52, 40), medial prefrontal cortex (mPFC; -1, 49, -5) as well as the right (Right angular gyrus; 46, -70, 36) and left (left angular gyrus; -46, -70, 36) angular gyri. As is standard in studies examining DMN connectivity, global connectivity of the PCC is reported. The remaining three ROIs made up the task-positive network and were centred in the intraparietal sulcus (IPS; -25, -57, 46), the frontal eye field region of the precentral sulcus (FEF; 25, -13, 50), and the middle temporal region (MT^{+}; -45, -69, 2). Seed-based connectivity results from the IPS are also reported to show DAN connectivity. Spherical ROIs were created using the MarsBaR toolbox (http://marsbar.sourceforge.net).

2.2.7. ROI-to-Voxel Functional Connectivity

We calculated ROI-to-Voxel (seed based) connectivity from one ROI for both the DMN (PCC seed) and DAN (IPS seed) for each behavioural condition: 1) Resting-state, 2) Block Motor Task, and 3) Event-related SST. Each analysis used two sets of the same functional data with the only difference between the sets being the temporal filtering (highpass vs. bandpass filtering).

In the SST, functional connectivity was measured for events when subjects successfully completed a go trial (go-success). Go-success trials were used because a) they occurred more often than the other events (i.e. go-failure, stop-success, stop-failure) and b) we are interested in the general effects of bandpass filtering on an event-related design, not on stop inhibition specifically.

Both highpass and bandpass filtered images for resting state, motor task, and SST were entered into the Conn toolbox (Whitfield-Gabrieli & Nieto-Castañón, 2012). The Conn Toolbox is a set of Matlab scripts for running functional connectivity analyses that utilises SPM routines. This method computes temporal

correlations between a ROI and all other brain voxels using a General Linear Model approach. Data from any subject moving more than half a voxel (> 1.5mm) during scanning was removed. A total of 22 datasets (for each temporal filtering group, N = 22 highpass; N = 22 bandpass) were included in rest and motor analysis. The SST included 21 datasets (N = 21 highpass; N = 21 bandpass). One subject was excluded from the SST due to excessive head motion during scanning.

Resting state, motor task, and SST ROI-to-voxel parameter estimate (Beta coefficient) maps for each ROI in both bandpass and highpass datasets were calculated for each subject and entered into second-level analyses using a one sample t-test in SPM. Age was entered as a covariate in all second-level analyses. We report results that survived $p < 0.001$ uncorrected at the voxel level and $p < 0.05$, FWE-corrected for multiple comparisons at the cluster level.

To identify statistical differences in connectivity maps for different filtering sets, resting state, motor task, and SST seed-to-voxel parameter estimate (Beta coefficient) images for each ROI in both bandpass and highpass datasets were calculated for each subject and entered into second-level analyses using paired t-tests in SPM. Age was entered as a covariate in these second-level analyses. We report results at voxel level $p < 0.001$ uncorrected and cluster level $p < 0.05$, FWE-corrected for multiple comparisons.

2.2.8. ROI-to-ROI Functional Connectivity

We also directly compared connectivity between the seven ROIs (four for DMN and three for DAN) in the three behavioural states. We constructed 7x7 matrices of Fisher z-transformed bivariate correlation coefficients (Pearson's r) using the ROIs described above. A hierarchical clustering analysis that places functionally similar ROIs in either DMN or DAN domains was conducted to assess whether any region changed its network membership across behavioural state. All connections between ROIs were FWE corrected ($p < 0.05$) at the ROI and network levels using the Network Based Statistics (NBS) toolbox. This method uses non-parametric permutation testing of network intensity to identify which ROI-to-ROI functional connections are statistically significant (Zalesky, Fornito, & Bullmore, 2010).

2.2.9. Relating DMN connectivity with reaction time

Connectivity within the DMN was measured and correlated with reaction time in the go success trials for the stop signal task with the aim of assessing the relationship between DMN connectivity and a measure of behavioural performance at the subject-level. This analysis involved averaging the edge strength (Pearson correlation coefficients normalised to z scores using the Fisher transformation) between PCC and the three DMN nodes (mPFC and bilateral angular gyrus). The final value is a measure of network-wide correlations for each subject in the analysis. This step was performed for both highpass and bandpass filtered sets, resulting in two different groups (Highpass PCC-DMN Connectivity, Bandpass PCC-DMN Connectivity). These values were then correlated (Pearson's Correlation using R studio (https://www.rstudio.com/)) with each subject's reaction time score. An equivalent analysis was not carried out for the motor task because performance was not assessed beyond ascertaining visually that participants were carrying out the task.

2.3. Results

2.3.1. Effects of aCompCor and temporal filtering on BOLD signal

The present study focused on understanding the role of high frequencies in DMN connectivity and their behavioural relevance. Figures 6 (A-F) show the power spectrum (FFT) from the average signal obtained from the grey matter mask of one subject used in the analysis. The top row panels show the raw data and the bottom row shows the data denoised with aCompCor, with the dotted line signifying the highpass cut-off. As expected, most of the power is contained within low frequencies (less than 0.08 Hz). However, it is clear that some power is contained within higher frequencies as well, suggesting that potentially relevant activity is removed from functional connectivity analyses that implement bandpass filtering. To control for possible confounds, we used the Nyquist frequency as a lowpass filter in the highpass set to be sure all the signals analysed were reliably sampled. The following analyses aim to understand the functional and behavioural role of these high frequencies in relation to DMN connectivity.

2.3.2. The effects of high frequency fluctuations in network dynamics

Functional connectivity between the PCC ROI and all other voxels in the brain during rest (Figure 7A), motor task (Figure 7B), and stop signal task (Figure 7C) was used to measure DMN functional connectivity in both highpass and bandpass sets. Across all three states, both temporal filtering sets show similar positive functional connectivity with DMN regions including PCC, mPFC, bilateral angular gyrus and medial temporal lobe structures. However, a qualitative difference was observed between temporal filtering groups when examining anticorrelations between DMN and DAN. Both sets showed anticorrelations in anterior regions of the DAN. However, the bandpass group showed greater anticorrelations with posterior DAN, most notably in bilateral inferior parietal cortex. This effect holds across all three behavioural states but is most prominent in the stop signal task. Functional connectivity between IPS and all other brain voxels was also measured for each behavioural state (Figures 8A-C). Connectivity between IPS and other DAN regions, including dorsolateral prefrontal cortex, anterior cingulate, and right IPS, was observed. Again, we found a qualitative increase in anticorrelations with DMN in bandpass filtered data compared to highpass filtered data while subjects performed a task.

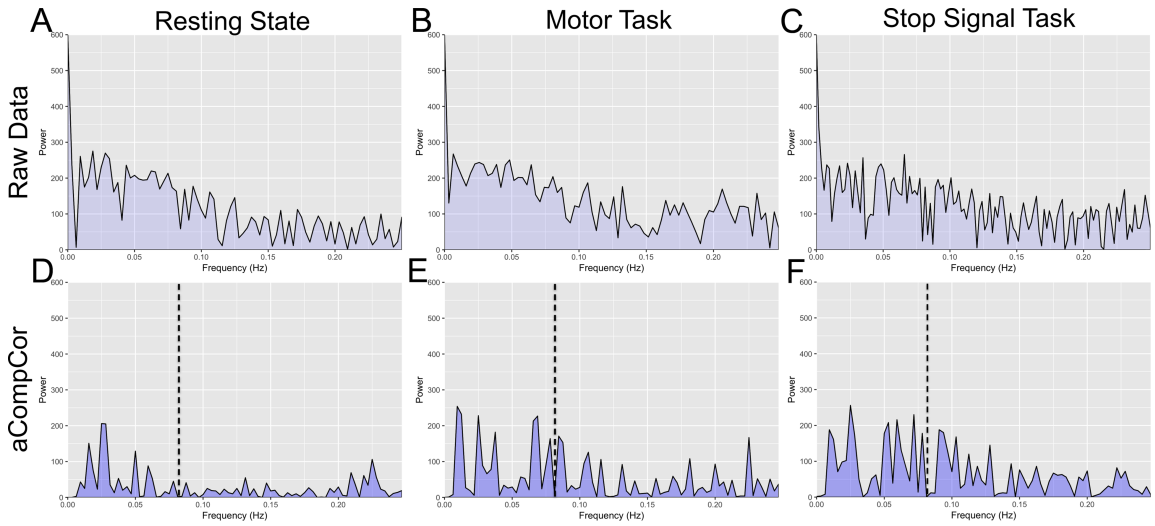

Figure 6. The FFT of the average grey matter signal from one participant
The top row (**A**-**C**) shows the FFT of raw signal, before any temporal pre-processing. The bottom row (**D**-F) shows data that was denoised with aCompCor. The dotted line (located at 0.08 Hz) indicates where the lowpass filter in the bandpass filter group is applied. The highpass group contains frequencies above (to the right of) this dotted line (up to 0.25 Hz, i.e. the Nyquist frequency).

2.3.3. Differences in functional connectivity between highpass and bandpass sets

To examine whether the differences in functional connectivity between temporal filtering sets observed in the PCC were statistically significant, we used paired t-tests to compare connectivity maps for each condition (Figures 9A-C). We also measured functional connectivity from IPS to all other voxels to determine whether this effect is observed in a putatively DAN (Figures 10A-C). For the highpass > bandpass contrast, highlighted regions (in blue) represent a lower magnitude of anticorrelation in the highpass group compared to the bandpass group. For the bandpass > highpass contrast, highlighted regions (in red) represent a greater magnitude in correlation in the bandpass group compared to the highpass group. All results are significant at voxel level $p < 0.001$ (cluster $p < 0.05$ FWE-corrected).

At rest, the PCC showed less anticorrelation in highpass compared to bandpass in many regions of the DAN (Figure 9A). These include bilateral insula, premotor cortex and inferior parietal cortex, middle frontal gyrus, superior frontal gyrus, right supplementary motor area, and several cerebellar regions. The PCC showed greater connectivity with DMN regions in the bandpass > highpass contrast, most notably the precuneus, mPFC, bilateral angular gyrus, and middle temporal gyrus. For the motor task (Figure 9B), the PCC followed the same pattern as rest, with greater anticorrelation in the bandpass set for highpass > bandpass, and greater correlation in the bandpass group in the bandpass > highpass contrast. Greater anticorrelation in bandpass filtered data was observed in the bilateral insula, supramarginal gyrus (SMG), right supplementary motor area, and right IFG. The bandpass > highpass contrast showed greater correlation between PCC and bilateral angular gyrus, mPFC, PCC, right IFG, and right ITG for bandpass filtered data. In the SST (Figure 9C), reduced anticorrelation was observed in the highpass group compared to the bandpass group in bilateral MFG, IPC, SFG, right MTG, left insula, as well as right and left IOC. Increased correlation was seen in the bandpass group in the bandpass > highpass group in DMN regions, including left angular gyrus, bilateral PCC, mPFC, and inferior frontal gyrus.

The IPS was less anticorrelated with several DMN regions in the highpass set (Figure 10A), including mPFC and right angular gyrus. There was also less

anticorrelation with bilateral superior temporal gyrus and several motor regions, including right supplementary motor area and right precentral gyrus. The IPS showed greater correlation with DAN regions in the bandpass set for the bandpass > highpass contrast. These regions include the left inferior occipital cortex (IOC), left inferior parietal gyrus, left inferior frontal cortex (IFG), left inferior parietal cortex (IPC), right precentral gyrus, and right cerebellar crus II. For the motor task, the IPS (Figure 10B) showed less anticorrelation in the highpass set with DMN regions including the mPFC, PCC, and right angular gyrus. In the bandpass > highpass contrast, greater correlation is seen in the bandpass group between IPS and left MTG, right IFG, bilateral SFG, left anterior cingulate (ACC), and left superior occipital cortex. In the SST, the IPS (Figure 10C) reduced connectivity in the highpass group in the highpass > bandpass contrast in several DMN regions including bilateral mPFC, precuneus, PCC, parahippocampal gyrus, and right ITG. For the bandpass > highpass contrast, IPS showed greater correlation in the bandpass group in DAN regions including left IFG, posterior ITG, superior parietal gyrus (SPG), MFG, IFG, and right precentral gyrus.

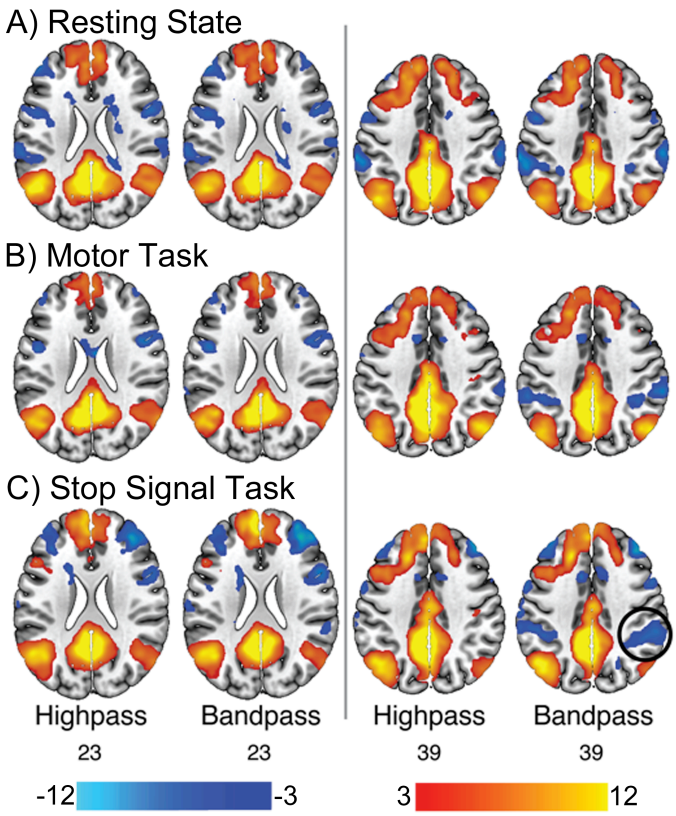

Figure 7. Whole brain connectivity analysis using PCC as a seed region
The differences between the two groups are most prominent in the right columns for the stop signal task (highlighted with a black circle). The first and third columns show highpass filtered data and the second and fourth columns show bandpass filtered data. **A** shows PCC connectivity during resting state, **B** shows PCC connectivity during the motor task and **C** shows PCC connectivity during the stop signal task. Red maps are positive correlations (t values) and blue maps are anticorrelations. Numbers below the axial slices are MNI z coordinates.

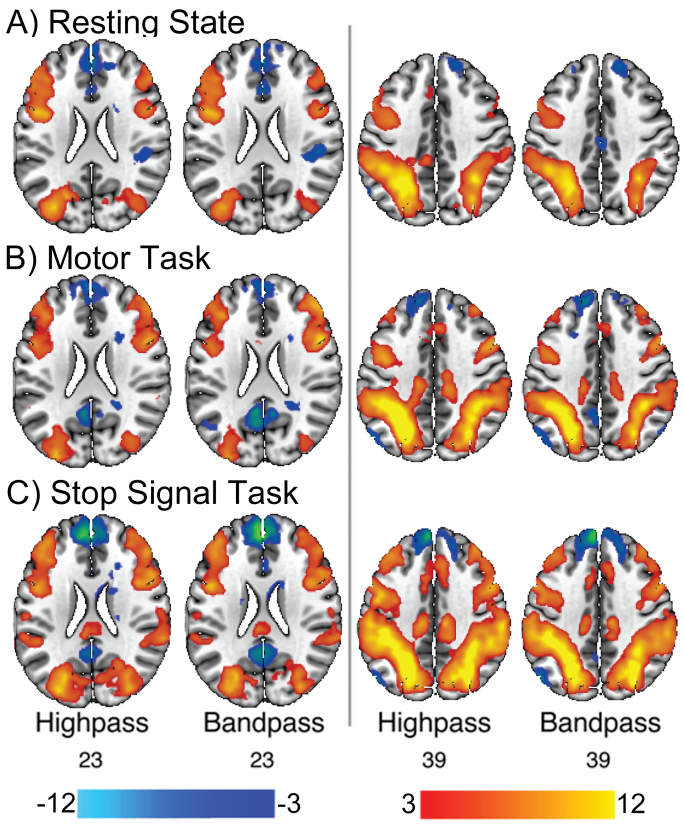

Figure 8. Whole brain connectivity analysis using IPS as a seed region
The first and third columns show highpass filtered data and the second and fourth columns show bandpass filtered data. **A** shows IPS connectivity during resting state, **B** shows IPS connectivity during the motor task and **C** shows IPS connectivity during the stop signal task. Red maps are positive correlations (t values) and blue maps are anticorrelations. Numbers below the axial slices are MNI z coordinates.

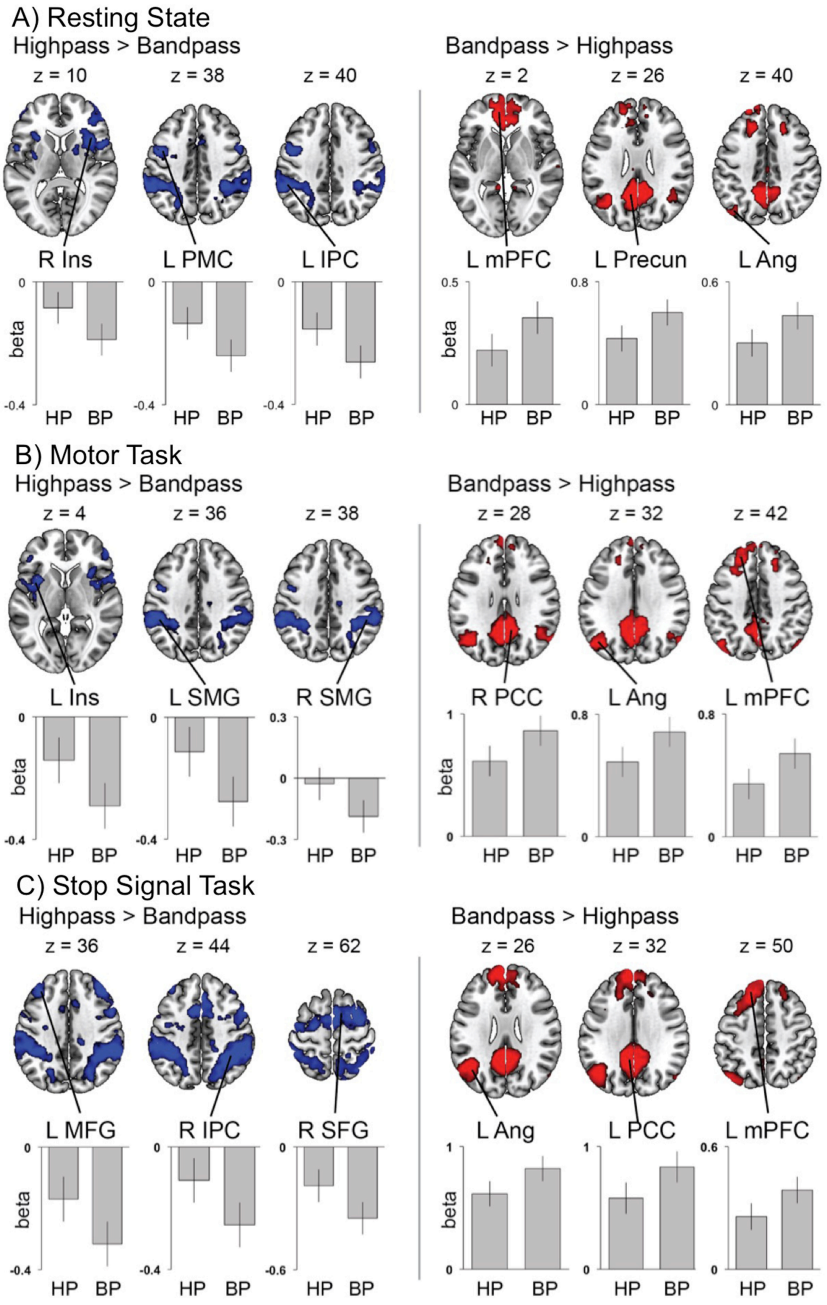

Figure 9. PCC functional connectivity between highpass and bandpass filtered data
Each subject's first level connectivity map for both filtering sets was entered into a paired t-test. All activations are statistically significant at $p < 0.001$ (cluster $p < 0.05$ FWE-corrected). The highpass greater than bandpass (blue regions) contrast is primarily located in TPN regions. The bandpass greater than highpass (red regions) contrast is primarily located in DMN regions. **A** shows resting state data, **B** shows Motor Task data and **C** shows SST data.

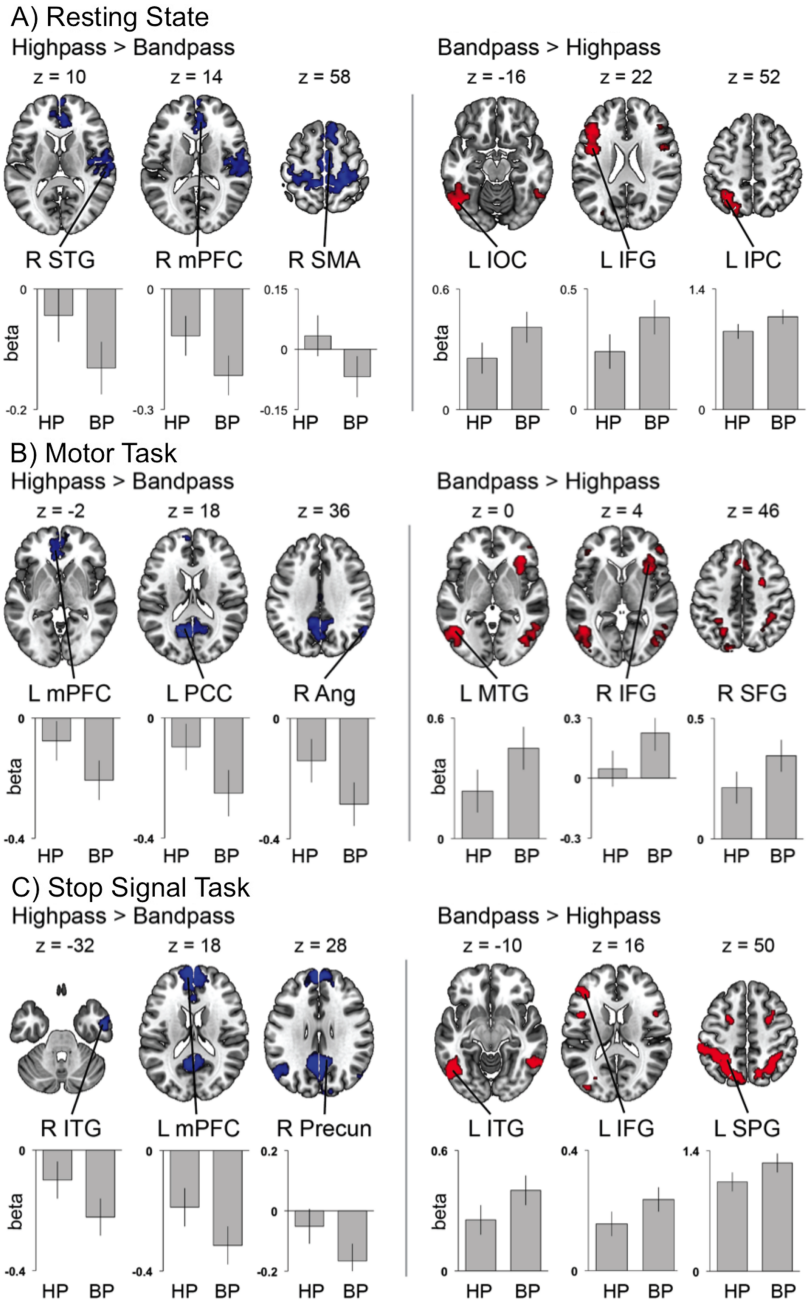

Figure 10. IPS functional connectivity between highpass and bandpass filtered data
Each subject's first level connectivity map for both filtering sets was entered into a paired t-test. All activations are statistically significant at $p < 0.001$ (cluster $p < 0.05$ FWE-corrected). The blue regions (highpass greater than bandpass) primarily shows reduced anticorrelations between IPS and DMN. The red regions (bandpass greater than highpass) primarily shows within TPN connectivity. **A** shows resting state data, **B** shows Motor Task data and **C** shows SST data.

2.3.4. Task-related internetwork connections are dependent on of high frequenices

To directly measure the effect of behavioural task and high frequencies on network reconfiguration, we measured functional connectivity between 7 previously defined ROIs representing regions of the DMN (PCC, mPFC and bilateral angular gyrus) and DAN (IPS, FEF, MT^+; Fox et al., 2005). Our results show that within network connectivity is relatively stable regardless of temporal filtering used, however between network positive connectivity is reliant on the inclusion of high frequencies in the data. These results are represented as chord diagrams in Figure 11A-F.

First, using a hierarchical clustering algorithm, we show that nodes from both DMN and DAN were reliably allocated to their respective networks for each behavioural task in both temporal filtering sets, suggesting within network connectivity is relatively stable across behavioural task and temporal filtering set. Second, anticorrelations between networks exist for both temporal filtering sets, but are much more prominent in the bandpass set, complementing our ROI-to-Voxel results from the previous section. These anticorrelations also increase during task state for both temporal filtering sets. However, a positive, long-range connection between mPFC in the DMN and MT^+ in the DAN is revealed only when high frequencies are included in the data. Crucially, this connection does not exist during resting state, and persists during both tasks, suggesting it is related to task engagement and not present due to high frequency noise, as that would likely be uniformly present across all three behavioural states. It is also not likely to have originated from residual motion-related artefacts, as those tend to result in increased short-range and bilateral connections and decreased long-range, anterior-posterior connections (Murphy et al., 2013; Power et al., 2012). We additionally examined connectivity during the fixation block of the motor task and found no internetwork positive connectivity (data not shown).

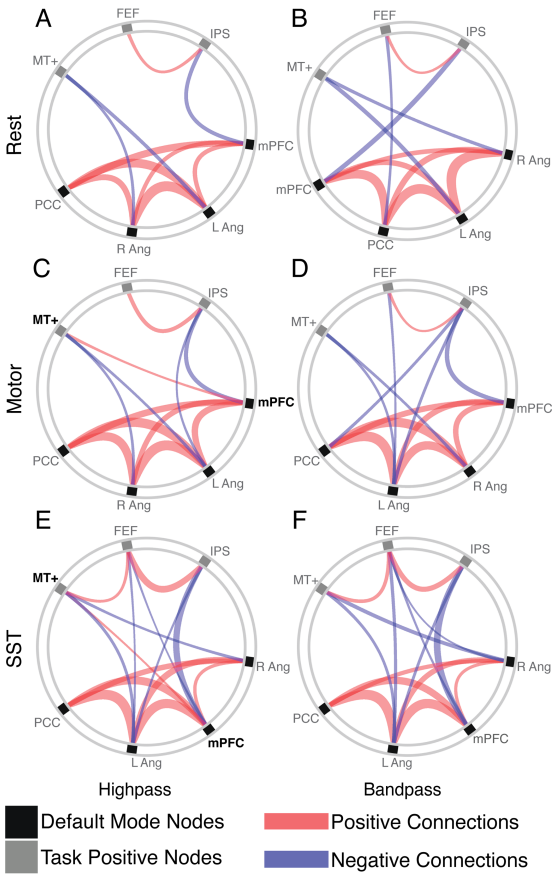

Figure 11. Connectivity patterns between DMN and TPN nodes
Chord diagrams showing ROI-to-ROI connectivity patterns between primary nodes in the DMN and TPN. **C** and **E** are of particular interest because they contain a positive connection between MT+ and mPFC, nodes (highlighted in bold) within the TPN and DMN respectively. All connections survived correction for multiple comparisons ($P < 0.05$, FWE-corrected) using the NBS toolbox (Zalesky etal., 2010). Node placement around the circle is determined by a hierarchical clustering algorithm, which reliably separated each node into their respective networks for each analysis group. Red lines are positive connections and blue lines are negative connections. The strength of connectivity is proportional to thet hickness of the line. Grey boxes show TPN nodes and black boxes show DMN Nodes. **A** shows connectivity in the highpass set at rest. **B** shows connectivity in the bandpass set at rest. **C** shows connectivity in the highpass set during execution of a finger opposition motor task. **D** shows connectivity in the bandpass set during execution of a finger opposition motor task. **E** shows connectivity in the highpass set during go success trials of a stop signal task. **F** shows connectivity in the bandpass set during go success trials of a stop signal task.

2.3.5. Inclusion of high frequencies reveals task-related DMN connectivity

In addition to examining network reconfiguration in the context of a task, we were also interested in assessing the influence of high frequencies in DMN on task-related behavioural measures. To this end, we correlated connectivity within the DMN (edge correlation strength of PCC to other DMN nodes) with reaction time in go success trials in the stop signal task for both highpass and band-pass filtered datasets. This analysis revealed a significant relationship with reaction time in the highpass group ($R^2 = 0.19$ r = -0.43, p = 0.046) (Figure 12A) but not the bandpass group ($R^2 = 0.084$ r = -0.29, p = 0.2) (Figure 12B), signifying that as within DMN connectivity increased, reaction time was faster. This could suggest that task-related connectivity exists exclusively in high frequency fluctuations in the DMN, and that bandpass filtering removes these behaviourally relevant signals.

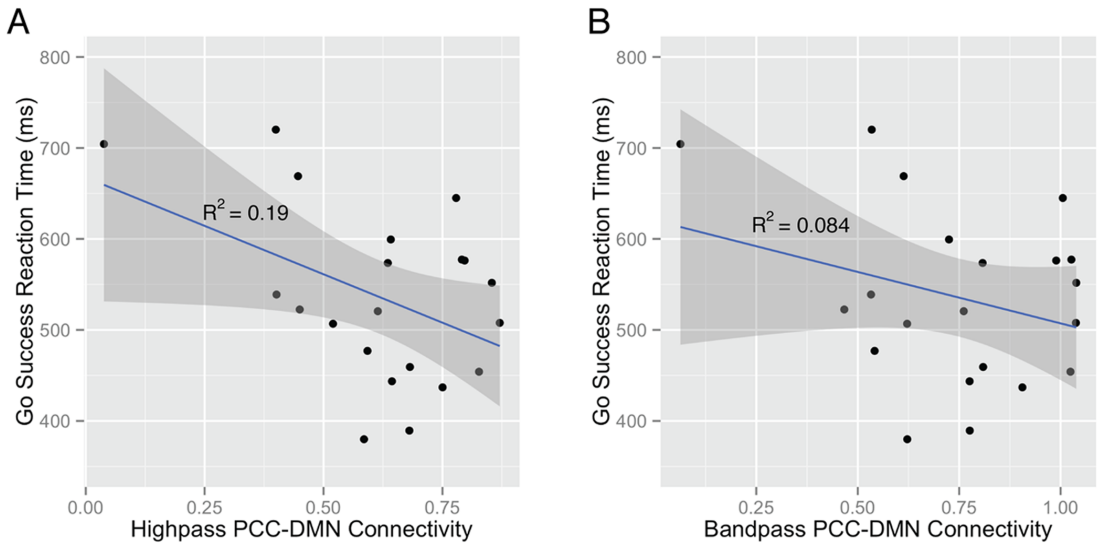

Figure 12. Pearson's Correlation of Go Success Reaction Time in the Stop Signal Task
A shows a significant negative correlation between highpass DMN connectivity with reaction time (R2 = 0.19, r = -0.438, p = 0.046), showing that reaction time is faster in subjects who have greater within-DMN connectivity. **B** shows a non-significant relationship between bandpass DMN connectivity and reaction time (R2 = 0.08, r = -0.29, p = 0.2).

2.4. Discussion

Several recent studies have focused on measuring the role of high frequencies in functional connectivity during resting state scanning (Boyacioglu et al., 2013; Chen & Glover, 2015; De Domenico et al., 2016; Gohel & Biswal, 2015; Kalcher et al., 2014; Lewis et al., 2016; Wu et al., 2008). At the same time, increasing evidence is suggesting that the DMN is not a "task-negative" network, but plays an active role in mediating a variety of cognitive processes (Spreng, 2012; Spreng, Sepulcre, Turner, Stevens, & Schacter, 2013; Vatansever et al., 2016, 2015a, 2015b, 2017). In an effort to bridge these two emerging lines of evidence, the aim of the present study was to identify whether inclusion of high frequency fluctuations in fMRI data alter DMN connectivity patterns and whether these alterations are task-relevant. To this end, we used fMRI data from three behavioural states (resting state, motor task, SST) processed with either a bandpass (0.009 Hz < f < 0.08 Hz) or highpass (0.009 Hz < f < 0.25 Hz) filter. DMN connectivity was measured and compared for each behavioural state and both filtering sets using ROI-to-Voxel and ROI-to-ROI connectivity. Using PCC as a seed region, we found that anticorrelations between PCC and posterior regions of the DAN were significantly reduced when high frequency fluctuations were included, an effect observed across all three behavioural states, though most prominent in the event-related SST. We also observed the inverse effect when measuring connectivity from IPS, a region of the DAN. Complementing our ROI-to-voxel result, we measured connectivity between DMN and DAN ROIs and observed stronger anticorrelations in the bandpass group across all three behavioural states, suggesting that between-network anticorrelations primarily fluctuate in low frequencies. Furthermore, we found that when high frequencies are included in the data, positive connectivity is observed between the mPFC and MT^{+}. Critically, this connection is only seen when subjects are engaged in a task and not during resting state. Our final analysis extended our initial findings to examine the behavioural relevance of within-DMN connectivity for individual subjects. Here, we found that greater within-DMN positive connectivity predicted faster reaction times in the highpass set, but not in the bandpass set. This shows that behaviourally relevant signals contained within the DMN exist at high frequencies, which are often removed

in functional connectivity analyses attempting to establish the DMN's role in behaviour.

We show that PCC anticorrelations are stronger in the bandpass set compared to the highpass set. Conversely, anticorrelations between IPS and regions of the DMN were also stronger in the bandpass set. Importantly we see this effect during all three behavioural states. These results show anticorrelations are greater at low frequencies when using seed regions from different networks, suggesting that this result is not specific to DMN regions. We also observed greater within DMN connectivity in the bandpass set compared to the highpass set. It has been suggested that anticorrelations serve to segregate neuronal processes that subserve opposite goals or competing representations (Fox et al., 2005). This relationship between DMN and DAN has been extensively explored, with many research groups suggesting a dynamic set of processes wherein DMN is involved in internally directed, self-generated cognition and DAN is involved in externally directed cognition (Andrews-Hanna, Smallwood, & Spreng, 2014). There are multiple task-positive networks, including the DAN, cingulo-opercular network (CON), and FPCN (Power et al., 2011). The DAN, in particular, has been shown to have an anticorrelated or antiphase relationship with DMN, while the FPCN is thought to play a mediating role between networks. Importantly, the majority of studies focusing on the relationship between DMN and various TPNs have applied bandpass filtering in the pre-processing stage, thus removing high frequencies from the statistical analysis. These anticorrelations have therefore been defined primarily in low frequencies. A likely interpretation could be that anticorrelations exist disproportionately but are not exclusively present at low frequencies and that bandpass filtering allows for a sharper focus on this phenomenon. This is supported by evidence suggesting that anticorrelations are present in large-scale networks independent of pre-processing techniques or research groups (Chai et al., 2012; Spreng, Stevens, Viviano, & Schacter, 2016).

In addition to finding that anticorrelations are behaviourally relevant and exist primarily in low frequencies, we provide two pieces of complementary evidence that suggest the DMN contains task-related activity at frequencies higher than what are typically analysed in functional connectivity studies. First, we show that when directly

comparing connectivity profiles between highpass and bandpass sets, a positive connection is observed between MT^+ and mPFC only during tasks, not during rest. This connection is unlikely to have originated from noise, as one would expect noise to be uniform across all three fMRI datasets. Furthermore, it is not likely to be related to head motion because motion-related artefacts have been shown to increase short-range and bilateral connectivity and decrease long-range and anterior-posterior connectivity (Murphy et al., 2013; Power et al., 2012). The positive connection we observe is both long-range and anterior-posterior. Second, we found that within-DMN connectivity reliably correlated with reaction time only in the highpass set. Previous studies have retained high frequencies when comparing task-related connectivity with resting state (Cole et al., 2014, 2016; Schultz & Cole, 2016), however, to our knowledge, this is the first demonstration that the DMN displays task-related connectivity only when high frequencies are included in the analysis. These results should be seen as a first step to elucidating network connectivity at higher frequency bands than usually considered as well as an investigation into the relationship between high frequency connectivity and behaviour. Further research with larger sample sizes and faster TRs (TR ≤ 0.5s) is necessary to draw more definitive conclusions.

Several recent studies have shown that activity in large-scale networks is not restricted to frequencies below 0.1 Hz. A previous study used a high sampling frequency (TR = 645 ms) to measure functional networks across 5 frequency bands at rest (Gohel & Biswal, 2015). Gohel and Biswal (2015) were able to identify well-established networks, including the DMN and several TPNs in high frequency bands. They also observed significant interhemispheric connectivity across all frequency bands, suggesting functional integration between brain regions is a multi-band phenomenon. Concurrently, recent work has shown the DMN to be involved in a variety of cognitive tasks involving both internally and externally directed attentional processes (Andrews-Hanna et al., 2014; Spreng et al., 2013; Vatansever et al., 2016, 2015a, 2015b, 2017). In this context, our work expands upon both lines of research by finding task-related activity in the DMN in high frequencies. This work may also find application in the field of mind-wandering and unconstrained cognition. Likely due to its role in internally directed attention, the DMN has frequently been implicated in mind-wandering, however, a recent meta-analysis showed that several studies have

reported connectivity between non-DMN regions during mind wandering (Fox, Spreng, Ellamil, Andrews-Hanna, & Christoff, 2015). Another recent study used a novel experience sampling method to show that DAN activity increased when subjects reported controlling their attention, while DMN activity increased when subjects reported a state of internal mentation (Van Calster, Argembeau, Salmon, Peters, & Majerus, 2017). These studies show that what has come to be known as resting state is a dynamic state where internetwork connectivity likely plays an important role. Therefore, mind-wandering and studies examining unconstrained cognition should consider preserving high frequency fluctuations in their data.

There are many remaining questions regarding high frequency fluctuations and the DMN. The present study is limited in that it only measures connectivity using *a priori* defined sets of ROIs. One approach could be to examine whether different subregions of the DMN communicate with other networks at different frequencies (Bzdok et al., 2015; Leech, Kamourieh, Beckmann, & Sharp, 2011). The use of data-driven approaches or multivariate analyses could extend this work. One such study by Crittenden et al. (2015) used multivoxel pattern analysis to show that DMN encodes task relevant information during performance of a task-switching behavioural paradigm. Though it was not the focus of the study, it should be mentioned that a low frequency bandpass filter was not applied to the data prior to statistical analysis. Dynamic connectivity methods are another important extension of this work. A recent study by Dixon et al. (2017) present a comprehensive analysis of DMN-DAN interactions across time as well as between different cognitive states and DMN subsystems (Dixon et al., 2017). Using dynamic functional connectivity to investigate the ever-changing nature of the relationship between DMN and DAN, the authors show that anticorrelation between these networks is transient rather than persistent and that anticorrelation is most prominent with core DMN regions like the PCC, mPFC, and angular gyri. As such, another interesting extension of our work could be the investigation of high frequency connectivity in different DMN subsystems.

Previous work looking to measure high frequencies in fMRI during resting state have used frequency bands originally defined by electrophysiology studies (Penttonen & Buzsáki, 2003) that argue frequency bands oscillate following a natural

logarithmic function. The current study did not use this definition of frequency bands because we were limited by the TR of 2s, meaning we could only reliably measure oscillations slower than the Nyquist frequency of 0.25 Hz. Future studies should use novel MR acquisition protocols, including multiplex and multiband EPI sequences (Boubela et al., 2013; Feinberg et al., 2010; Kalcher et al., 2014), that allow for faster TRs, and thus higher Nyquist frequencies, to further define the role of high frequencies during rest and task based functional connectivity. Using these acquisition protocols would also reduce the potential impact of physiological and motion-related artefacts. A potential area of controversy in the functional connectivity literature that applies to the current study is the use of denoising pipelines. We acknowledge that there are a variety of potential denoising methods that can influence functional connectivity results, including manual denoising using ICA (e.g. fsl_regfilt) and automated denoising methods that implement PCA (Kay et al., 2013) or ICA (Salimi-Khorshidi et al., 2014). We chose to use aCompCor because it has been shown to sufficiently remove noise when studying network interactions (Chai et al., 2012). However, a replication of these findings by researchers using alternative denoising pipelines is probably warranted before definitive conclusions on this matter are drawn.

Two main conclusions that can be drawn from this study. First, we found that the inclusion of high frequencies in functional connectivity analyses results in differences in between network interactions. Second, this high frequency activity appears to be important particularly when examining connectivity during task performance. This is supported by two sets of results, both of which delineate the role of high frequencies in the DMN. First, we found differences in connectivity patterns while subjects perform an attentionally demanding task and second, DMN connectivity correlated with reaction time only when high frequencies are included. These results suggest that functional connectivity studies should also include high frequencies. This is especially the case when examining DMN connectivity during task execution. Therefore, the remaining experiments in this thesis all use highpass filtering during pre-processing.

Chapter 3. Propofol-induced alterations in brain connectivity reflect parvalbumin interneuron distribution in the human cortex

3.1. Introduction

General anaesthesia is a drug-induced, reversible state of unconsciousness that is maintained for the duration of surgical and other nonsurgical medical procedures. Anaesthesia results in specific physiological and behavioural characteristics including analgesia, amnesia, akinesia and physiological stability (Brown et al., 2010). One of the most frequently used anaesthetics is propofol, a GABAergic drug that increases inhibitory input to neurons throughout the central nervous system (Brown et al., 2011; Jurd et al., 2002; Rudolph & Antkowiak, 2004). Despite an understanding of its primary receptor target (Yip et al., 2013), little is known about how these molecular and cellular processes influence the macro scale cortical network mechanisms that result in propofol-induced anaesthesia.

Broadly speaking, cortical networks are composed of glutamateric excitatory projection neurons or principal cells (PC) that propagate signals within and between regions, and local GABAergic inhibitory interneurons that gate signal flow to generate cortical rhythms (Bartos, Vida, & Jonas, 2007; Tremblay, Lee, & Rudy, 2016). The population of GABAergic interneurons in the mammalian cortex can be divided into three categories, based on three largely non-overlapping markers; Parvalbumin (PV), Somatostatin (Sst) and 5-HT3aR (5HT3a) (Lee et al., 2010; Rudy, Fishell, Lee, & Hjerling-leffler, 2013; Tremblay et al., 2016). Parvalbumin interneurons represent 40% of cortical interneurons and include Basket cells, which target the soma and proximal dendrites of PCs, and Chandelier cells, which target the axon of PCs (Kawaguchi & Kubota, 1997). Somatostatin interneurons comprise approximately 30% of the cortical interneurons and are made up of Martinotti and non-Martinotti cells. Both of these subtypes target the dendrites pyramidal cells (Kubota, 2014). Finally, around 30% of cortical interneurons are marked by 5HT3a, an ionotropic serotonin receptor. These can be further subdivided into Vip (~40% of 5HT3a interneurons), which primarily

target other interneurons (Jiang et al., 2015) and non-Vip cells, which are involved in synaptic and volume transmission (Tremblay et al., 2016)

Multiple studies have shown that oscillatory brain activity is highly dependent upon a network of inhibitory GABAergic interneurons (Huang et al., 2016; Sohal, Zhang, Yizhar, & Deisseroth, 2009; Traub, Jefferys, & Whittington, 1997; Whittington & Traub, 2003; Whittington, Traub, & Jefferys, 1995). Sohal et al. (2009) used optogenetics in a rodent model to selectively modulate neuronal circuitry and found that driving PV interneuron firing resulted in increased gamma oscillations *in vivo*, while selective inhibition of PV interneurons suppressed gamma frequency rhythms demonstrating that GABAergic interneurons are important in driving synchronous brain activity and in turn proposing that disruption of these mechanisms by GABAergic anaesthetic drugs like propofol could result in loss of connectivity.

As previously reviewed, network dynamics during induction and recovery from propofol sedation have previously been studied using multiple neuroimaging modalities, including fMRI (Boveroux et al., 2010; Schrouff et al., 2011; Stamatakis et al., 2010), EEG (Chennu, O'Connor, Adapa, Menon, & Bekinschtein, 2016; Purdon et al., 2013; Supp, Siegel, Hipp, & Engel, 2011) and intracranial electrocorticography (ECoG; (Lewis, Weiner, Mukamel, Donoghue, & Eskandar, 2012). EEG shows an increase in β oscillations (15-30 Hz) during sedation induction with propofol (Purdon et al., 2013). This is followed by an increase of global slow (0.1 – 1.0 Hz) and δ (1-4 Hz) oscillations, as well as frontal α oscillations (8-12 Hz) when subjects are fully anaesthetised (Lewis et al., 2012; Mhuircheartaigh, Warnaby, Rogers, Jbabdi, & Tracey, 2013; Purdon et al., 2013).

Functional connectivity of the fMRI BOLD signal is a commonly used technique to study brain network dynamics (Biswal et al., 1995; Fox & Raichle, 2007). When examining functional connectivity during propofol sedation, Boveroux et al., (2010) found decreased connectivity within networks typically thought to be important for high-level cognition, including the DMN and the executive control network (ECN). These networks have been shown to be anticorrelated during resting wakefulness (Fox et al., 2005), however, Boveroux et al., (2010) found that as sedation

increased, the anticorrelation between these networks decreased. This body of work suggests that propofol induces a loss of consciousness by disrupting synchronous and asynchronous patterns within large-scale brain networks.

Despite the advances in understanding anaesthetic actions at the molecular and systems level, little work has focused on addressing the multi-scale mechanisms involved in propofol-induced sedation. With this in mind, the present study uses data-driven statistical methods with human functional imaging data to identify whether patterns of cortical network dynamics are reflected in the distribution of one or more GABAergic interneuron subtypes. First, cortical gene expression maps were generated for PV, Sst and 5HT3a receptors using the Allen Institute for Brain Science (AIBS) Human Microarray dataset. Next, we used fMRI data from 25 healthy participants undergoing propofol sedation. We then created functional connectivity matrices for each of the four sedation conditions and used a data-driven method, partial least squares (PLS), to identify patterns of connections that were collectively affected by propofol sedation. Finally, we correlated the cortical gene expression maps with regional connectivity profiles during propofol sedation.

3.2. Methods

3.2.1. Participants

Ethics for this study were obtained from the Cambridgeshire 2 Regional Ethics committee. 25 healthy participants were recruited for the study (9 female; mean age = 35; range = 19-35 years). Two senior anaesthetists were present during the entire session to observe the healthy volunteers.

3.2.2. Sedation Procedure

Propofol administration was delivered intravenously as a target controlled infusion (plasma concentration mode), using an Alaris PK infusion pump (Carefusion, Basingstoke, UK) programmed with the Marsh pharmacokinetic model (Absalom, Mani, Smet, & Struys, 2009; Marsh, White, Morton, & Kenny, 1991). This system automatically determines the required infusion rate to achieve and maintain a target

plasma concentration set by the anaesthesiologist. This model is routinely used in the clinic to control propofol infusion during general anaesthesia and sedation. In the present study there were four conditions, with three different plasma target levels: 1) no drug 1 (baseline), 2) 0.6 ug/ml (low sedation), 1.2 ug/ml (moderate sedation) and 4) no drug 2 (recovery). There was a 10 minute period that allowed for plasma equilibrium and effect-site propofol concentration before cognitive tests began. Blood samples were drawn (a total of 6) at the end of each titration period and before the plasma target was altered. Sedation level was assessed verbally before and after each scanning run. TIVATrainer (pharmacokinetic simulation software package; www.eurosiva.org) showed that plasma concentrations would asymptotically approach zero 15 minutes after propofol sedation. Therefore recovery scans were acquired at 20 minutes following cessation of propofol. The mean plasma propofol concentration was 304.8 (SD = 14.1) ng/ml during light sedation, 723.3 (SD = 320.5) ng/ml during moderate sedation and 275.8 (SD = 75.42) ng/ml during recovery. The mean total mass of propofol administered was 210.15 (SD = 33.17) mg, equal to 3.0 (SD = 0.47) mg/kg.

3.2.3. Data Acquisition

A Siemens Trio Tim 3T MRI system (Erlangen, Germany) was used to collect the data. A 12-channel head matrix transmit-receive coil, using a fast-sparse 32 slice axial oblique sequence was utilised (TR = 2s, TE = 30ms, flip angle = 78, voxel size = 3.0 x 3.0 x 3.0mm, matrix size 64 x 64, field of view 192mm x 192mm, slice thickness = 3.0mm, 0.8mm gap between slice, bandwidth = 2422 Hz/Px). 150 volumes were collected per sedation condition (600 across the entire scanning session). T1-weighted MPRAGE high-resolution structural images were also acquired at 1mm isotropic resolution (TR = 2250ms, TI = 900ms, TE = 2.99ms, flip angle = 9).

3.2.4. Pre-processing

Image pre-processing was performed using Statistical Parametric Mapping (SPM) 12 (http://www.fil.ion.ucl.ac.uk/spm/) and MATLAB Version 2016a (http://www.mathworks.co.uk/products/matlab/). The first 5 images for each sedation level were removed to eliminate saturation effects and achieve steady-state

magnetization. All functional data were slice-time corrected and motion corrected. Structural images were coregistered to the mean EPI and segmented into grey matter, white matter and cerebrospinal fluid masks, then normalised to MNI space with a resolution of $2mm^3$ (Ashburner & Friston, 2005). Functional images were then smoothed with a 6mm FWHM Gaussian kernel. To mitigate movement-related and physiological noise, data underwent de-spiking with a hyperbolic tangent squashing function, followed by the aCompcor technique. This process removes the first 5 principal components of the signal from white matter and CSF masks, as well as 6 motion parameters and their first-order temporal derivatives and a linear detrending term (Behzadi et al., 2007). Images were then highpass filtered to remove low frequency fluctuations attributable to scanner noise (0.009 Hz < *f*).

3.2.5. Creating functional connectivity matrices

The Conn functional connectivity toolbox was used to extract time series from a set of predefined regions of interest (ROI) (Whitfield-Gabrieli & Nieto-Castanon, 2012). The Lausanne parcellation, a commonly used brain region parcellation template, (Hagmann et al., 2008) was used to define a set of 129 cortical brain regions. Because the present study was focused exclusively on large-scale cortical brain networks, 11 subcortical brain regions were removed from the analysis, leaving a total of 118 regions. For each subject, the BOLD time series was divided into the four sedation conditions. ROI timecourses were created by extracting the signal in each volume for each brain region across subjects. Regional time series were then correlated, creating individual subject connectivity matrices for each sedation condition.

3.2.3. Partial Least Squares

Partial least squares (PLS) is a multivariate statistical analysis used to identify relationships between sets of variables (Abdi, 2010; Krishnan, Williams, Randal, & Abdi, 2011; McIntosh & Lobaugh, 2004; McIntosh & Mišić, 2013; Mišić et al., 2016). The analysis aims to identify linear combinations of variables in each set that maximally covary with each other. In the current analysis, this corresponds to a weighted combination of functional connections that maximally explain a pattern associated with changing levels of sedation.

3.2.3.1. Singular Value Decomposition

Functional data for each sedation condition were organised in a data matrix containing each unique connection (corresponding to the upper elements of the connectivity matrix) for each subject stacked on top of each other. With $n = 118$ nodes, this corresponds to $k = 118 \times (118 - 1)/2 = 6903$ unique connections per condition. Thus, for the Cambridge dataset, the data matrix contains $k = 6903$ columns and $N = 100$ rows (that is, 25 subjects x 4 conditions). We used a mean-centred approach where connection strengths are expressed as a deviation from the grand mean across the entire data set. This matrix, **X**, then underwent singular value decomposition (SVD; Eckart & Young, 1936):

$$\mathbf{X} = \mathbf{U}\Delta\mathbf{V}'$$

Such that

$$\mathbf{U}'\mathrm{U} = \mathbf{V}'\mathrm{V} = \mathbf{I}$$

This analysis outputs a set of orthogonal latent variables (LVs), where **U** and **V** are matrices of left and right singular vectors and Δ is a matrix with singular values along the diagonal and zeros elsewhere. Singular vector elements weigh the contribution of individual variables (connections across sedation condition) to the overall multivariate pattern. The proportion of covariance between functional connectivity and study design captured by an LV serves as a measure of effect size and is estimated as the ratio of the squared singular value to the sum of squared singular values from the SVD (McIntosh & Lobaugh, 2004).

3.2.3.2. Assessing significance of LVs with permutation tests

Each LV was tested for statistical significance using permutation tests, by randomly reordering the rows (i.e. condition labels: baseline, low sedation, moderate sedation and recovery) of the original **X** matrix. The permuted matrices were then mean centred and subjected to SVD as described above. This process creates a distribution of singular values under the null hypothesis so that there is no relationship between study design (sedation conditions) and functional connectivity. Since singular values

are proportional to the effect size of the LV, a p-value for each LV was estimated as the number of times the permuted singular values exceeded the original singular value.

3.2.3.3. Assessing the contribution and reliability of specific connections

We used bootstrap resampling to assess the contribution and reliability of each connection to the overall multivariate pattern. This process allows the patterns derived for each LV to be cross-validated. The algorithm creates new data matrices that randomly resamples participants with replacement and performs SVD generating a sampling distribution for each weight in the singular vector. This bootstrap distribution can then be used to estimate the standard error of each connection weight. We then calculated a "bootstrap ratio" for each connection by dividing the weight from the singular vector by its bootstrap-estimated standard error. A bootstrap ratio with a large magnitude (positive or negative) signifies that the connection (a) makes a large contribution to the LV and (b) is stable across participants.

3.2.4. Brain network analysis

The output of our PLS analysis consisted of a single 118 x 118 matrix that contained a bootstrap ratio for each connection. When a bootstrap distribution is approximately normal it is equivalent to a z-score (Efron & Tibshirani, 1986). Thus, connections were thresholded at the 95% confidence interval (corresponding to z = 1.96).

Each ROI in our parcellation was overlapped with previously defined large-scale functional connectivity networks (Smith et al., 2009). We modified a subset of these networks by collapsing all visual networks into one and also combining the left and right frontoparietal network into one network. This resulted in 6 large-scale networks, that included the Auditory Network (Aud), DMN, FPCN, SN, Somatomotor Network (SM) and Visual Network (Vis). To assess the contribution of each network to the significant LV (i.e. the pattern of connectivity associated with changes in propofol sedation), we averaged the bootstrap ratios within and between each large-scale network, including any connections with a value of 0. This allowed us to control for the fact that some networks have more ROIs and therefore more possible connections than others.

3.2.5. Gene expression maps

The AIBS Human Brain Atlas is a publicly available resource of microarray gene expression profiles across a set of brain regions (Hawrylycz et al., 2012). A DNA microarray consists of a number of DNA probes attached to a fixed surface. It allows for the concentration of a large number of individual genes in a sample tissue to be quantified. The atlas uses post-mortem tissue from 6 donors, none of who had a history of neurological or neuropsychiatric disease and who passed a strict set of serology, toxicology and RNA quality screens. The donors were a 24-year-old, left-handed African American male (H0351.2001), a 39-year-old, left-handed African American male (H0351.2002), a 57-year old, ambidextrous Caucasian male (H0351.1009), a 31-year old, right-handed Caucasian male (H0351.1012), a 49-year old, right-handed Hispanic female (H0351.1015) and a 55-year old, right-handed Caucasian male (H0351.1016). Two subjects (H0351.2001 and H0351.2002) had data from both hemispheres collected. The remaining subjects only had data from the left hemisphere collected. Each brain was sliced into 0.5-1cm thick slices and frozen. Brains were then sectioned into 0.5-1.0 cm slices. Each section was then further divided to roughly represent each structure throughout the brain. RNA was isolated from each sample and microarray analysis was performed to produce expression levels in 60,000 genes. 93% of genes are represented by at least two probes (see technical white paper at *http://human.brain-map.org*)

We overlapped the centres of each region in our MRI parcellated data to the MNI coordinates of the samples from the AIBS dataset. If multiple probes were present in one ROI, we averaged the gene expression. We used the gene list from Whitaker and colleagues to discard probes that did not match gene symbols presented in the AIBS dataset (Whitaker et al., 2016). This generated a matrix with 20,737 rows (corresponding to different genes) and 118 columns (corresponding to different regions). The cells in this matrix indicate the expression of a particular gene in each brain region. To create gene expression maps for PV, Sst and 5HT3a we indexed this matrix taking all probes that corresponded to each gene. This gave us a column vector that represented concentration for each gene in each brain region.

Next, we used the unthresholded bootstrap ratio matrix and summed the values across each row, resulting in a column vector representing the contribution weight for each region. We then used a Pearson's correlation to assess whether there was a statistical relationship between a region's change in connectivity during propofol sedation and its expression of PV, Sst or 5HT3a. We used the unthresholded bootstrap ratio matrix because the thresholded matrix had many ROIs with no significant connections, and therefore had a number of values along the Y axis. All correlations with gene expression maps were significant using either the thresholded or unthresholded bootstrap ratio matrices.

3.3. Results

3.3.1. Distribution of GABAergic interneuron gene expression in the cortex

Figure 13 shows the concentration of each of the 3 genes of interest across the cortex. PV (Figure 13A) had a higher concentration in comparison to either SSTR or HTR3A. Highest concentrations were observed in the Vis, SM and midline areas associated with the DMN. SSTR (Figure 13B) was less concentrated than PV across the cortex, with high expression in higher cognitive networks including the FPCN, SN and DMN. Interestingly, SSTR showed a somewhat asymmetrical expression between the hemispheres, with qualitatively greater expression in left hemisphere compared to right hemisphere. HTR3A (Figure 13C) was the least concentrated of the three genes and was expressed mostly in Aud, DMN and SN.

Previous work has shown that gene expression between hemispheres is largely symmetrical (Chen et al., 2011), with the similarity being particularly pronounced in visual regions. Therefore, to validate the expression maps we created, we examined whether there were interhemispheric differences in expression of each of the GABAergic receptor subtypes (Figure 14). Specifically, we used t-tests to compare expression measures from the left and right hemispheres of the visual network (8 regions per hemisphere). We found no differences in the left, right comparisons of expression of PV ($p = 0.35$), SST ($p = 0.83$) or 5-HT3a ($p = 0.07$).

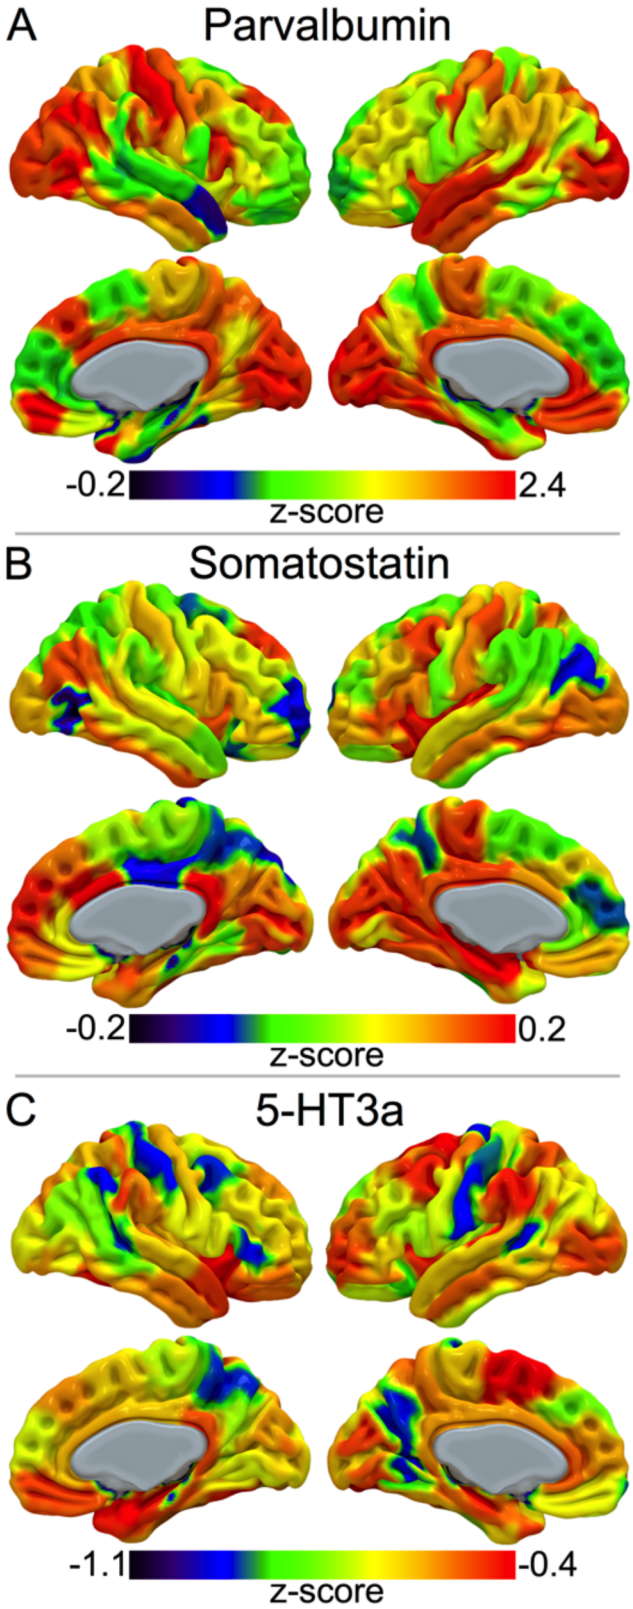

Figure 13. Cortical expression maps of GABAergic interneurons
Each map was created using the AIBS Human Microarray database. This database has microarray expression values for 20,737 genes across 6 subjects. The maps show an average value in each region across each of the 6 subjects **A** shows the cortical expression of parvalbumin. **B** shows the cortical expression of somatostatin. **C** shows the cortical expression of 5-HT3a.

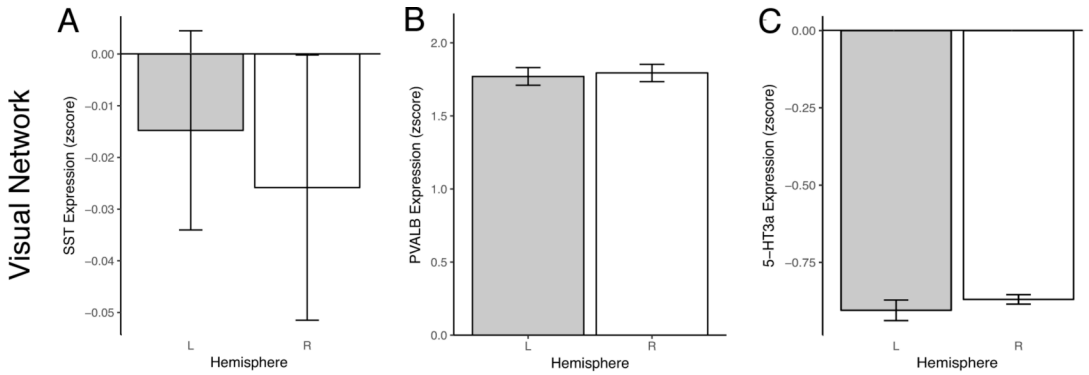

Figure 14. GABAergic interneuron gene expression between hemispheres
A, **B** and **C** show that there are no significant differences in expression of SST, PVALB or 5-HT3a between left and right hemispheres in the visual network. These measurements represent averages across each of the 6 subjects from the AIBS dataset.

3.3.2. Cortical functional connectivity patterns during propofol sedation

Next, we used partial least squares, a multivariate statistical analysis, to investigate cortical functional connectivity patterns during propofol sedation. This analysis served two purposes: 1) It determined which brain regions were most affected by propofol sedation, thus allowing us to correlate changes in connectivity with concentrations of different markers for GABAergic interneuron subtypes. 2) It quantified the dynamics of each within and between-brain network interaction, allowing us to compare previous findings in this field with our results using a completely data-driven method. PLS has not been used to study network dynamics during sedation before and we selected this technique because it encompasses the simultaneous observation and analysis of more than one outcome variables in contrast to seed-based functional connectivity methods.

The PLS technique decomposes the covariance between different sets of variables (McIntosh and Mišić 2013), in this case, the cortical functional connectivity pattern during four different sedation stages. Essentially, PLS characterises connectivity dynamics for each brain region without assuming specific relationships between them. The output of this PLS analysis was a set of 4 latent variables (LV), each with a corresponding set of weights for each connection. Each LV represents a weighted pattern of functional connectivity that optimally covaries with the changing

level of sedation. To determine statistical significance, each LV was subjected to permutation testing. This process identified only 1 of the 4 LVs and its associated weights to be significant (covariance accounted for = 22.3%, p = 0.000009), therefore this LV was selected for further analysis.

The resulting set of connections from this LV (Figure 15A) followed one of two inverse patterns of connectivity (Figure 15B). Each connection in the matrix corresponds to a bootstrap ratio (BSR), that is, the weight of a connection divided by its bootstrap estimated standard error. Using this metric is advantageous because it reveals connections that are both high in magnitude and consistent across participants. Connections with positive BSR scores follow a pattern of connectivity where, during the awake condition, connections displayed low levels of connectivity but as propofol concentrations increased, these connections increased connectivity, then during recovery from sedation, levels of connectivity dropped to near awake levels Brain regions that closely followed this pattern can be seen in Figure 15C. In contrast, connections with negative BSR scores were highly connected during the awake condition, then, as propofol concentrations increased, they decreased connectivity. During recovery, these connections then regained the level of connectivity they had during the awake condition. Brain regions that had connections following this pattern can be observed Figure in 15D. Patterns of functional connectivity, therefore, follow 1 of 2 inverse trajectories during propofol administration, with some connections increasing as a function of sedation while others decrease.

As the alteration in dynamics of large scale network connectivity has repeatedly been associated with consciousness alterations in both anaesthesia and patients with disorders of consciousness (Di Perri et al., 2016; Stamatakis et al., 2010), we next examined connectivity patterns within 6 canonical large-scale brain networks (Aud, DMN, FPCN, SN, SM and Vis). To do so, we constructed a 6 x 6 brain network matrix by averaging the positive and negative BSR scores within and between each network, reflecting the degree to which each network follows the positive or negative pattern

The Positive BSR matrix (Figure 16A) shows that DMN (Figure 16B) and FPCN both contain connections that increase in connectivity strength during sedation. These networks also show a significant number of between network connections that increase in strength as propofol sedation increases (Figure 16C). Additionally, we found that connections between DMN and Aud and SN also increase connectivity during propofol sedation. For the Negative BSR matrix (Figure 16D) we found that a subset of within-DMN (Figure 16E) and Vis-SM internetwork (Figure 16F) connections decreased connectivity during propofol sedation.

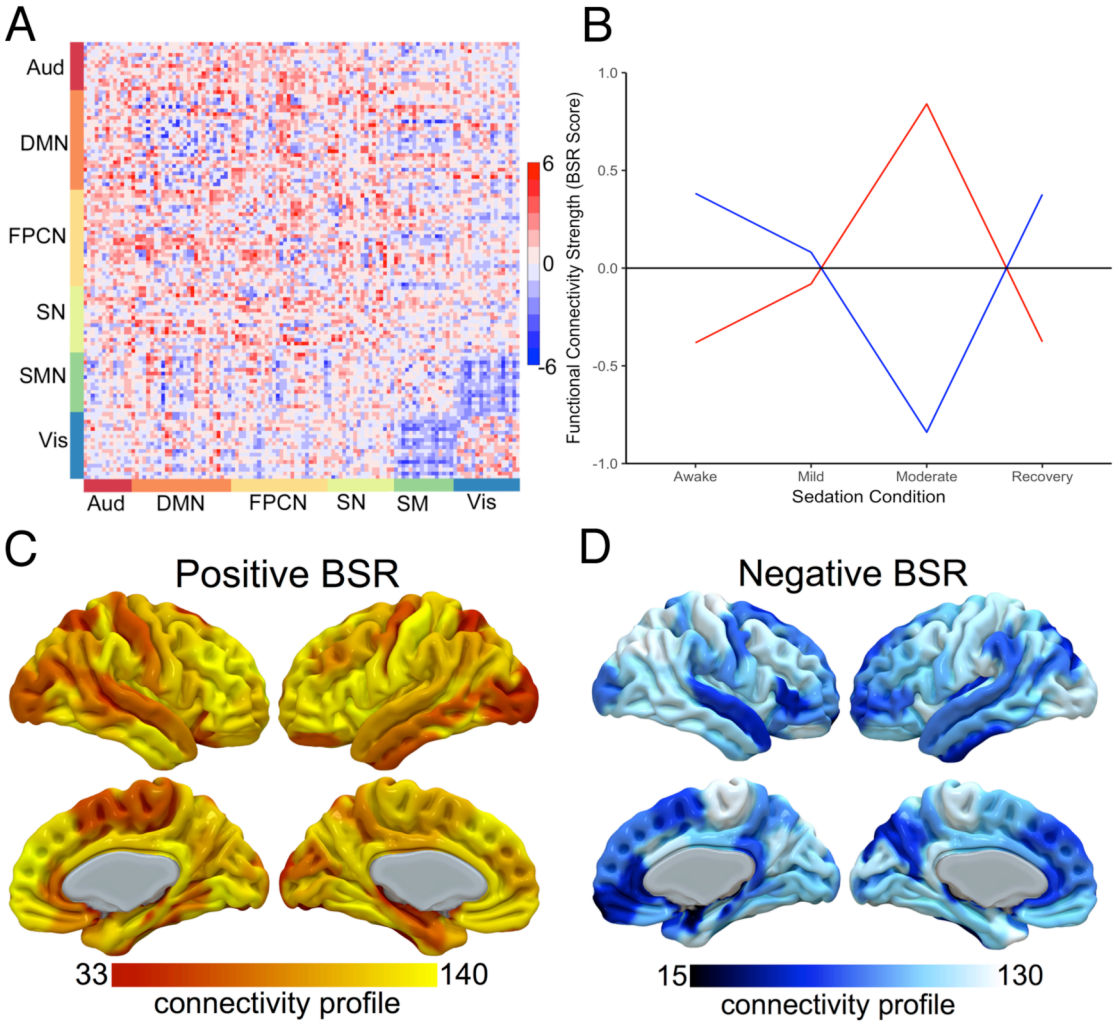

Figure 15. Inverse patterns of functional connectivity during propofol sedation
A shows the BSR scores for each connection between regions parcellated. These are organised into 6 large-scale cortical functional networks (Aud, DMN, FPCN, SN (SN), SM and Vis). **B** shows patterns of connectivity that corresponds to connections with positive BSR scores (red) negative BSR scores (blue). **C** and **D** show regional connectivity maps created by summing the BSR scores for each connection, resulting in an overall connectivity weight for each region. Higher values correspond to that region having a high number and highly weighted BSR scores that either gain or lose connectivity during sedation. **C** shows regions that gain connections during sedation. **D** shows regions that lose connections during sedation.

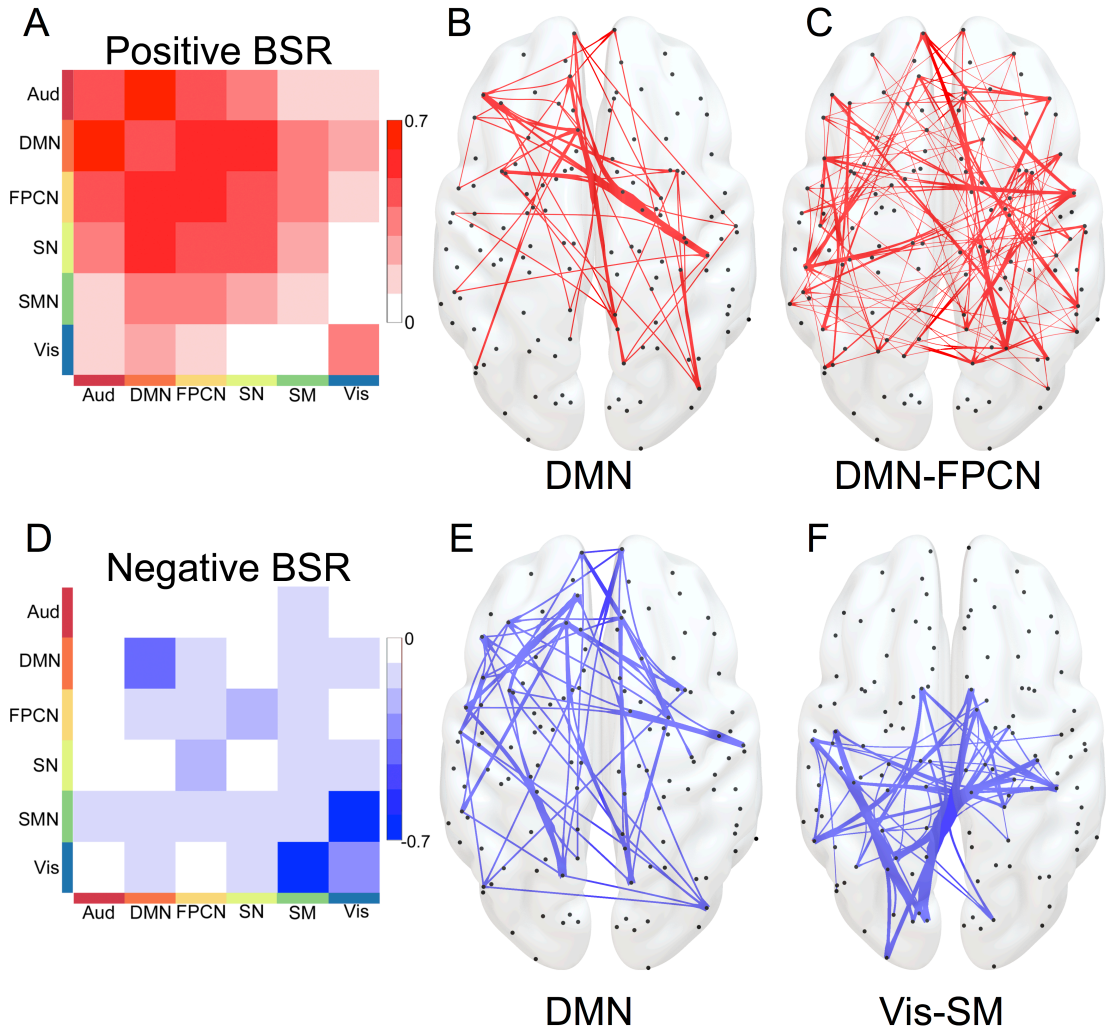

Figure 16. Network level alterations in functional connectivity during propofol sedation
A shows positive BSRs collapsed across large-scale networks and represents the degree to which individual networks increase within or between network connectivity during propofol sedation. **B** shows within-DMN and **C** shows DMN-FPCN internetwork connections that increase connectivity during sedation. **D** shows negative BSRs collapsed across large-scale networks and represents the degree to which each network decreases within and between network connectivity during sedation. **E** shows within-DMN and **F** shows Vis-SM internetwork connections during propofol sedation.

3.3.3. Propofol-induced loss of connectivity correlates with parvalbumin expression

Our final aim was to determine whether this changing pattern of functional connectivity is reflected by the distribution of a specific GABAergic interneuron subtype. To construct a connectivity profile for each region of interest we first split the BSR matrix into positive and negative values. Next, we summed each row of the matrix, giving us a value corresponding to how strongly that region followed either the positive or negative connectivity pattern. Figure 17A shows the positive regional connectivity profile, while Figure 17B shows the negative regional connectivity profile. For clarity, both strongly positive and negative correlations between regional time courses are considered functionally connected, while correlation values around 0 are considered disconnected.

We then correlated the regional connectivity profile for positive and negative BSR scores with the regional concentration of each gene of interest. A significant correlation here would suggest that there is a relationship between the connectivity profile and the presence of a GABAergic interneuron subpopulation. No significant relationship was found between SSTR expression and positive ($r = 0.093$, $p = 0.31$; Figure 18A) or negative ($r = -0.135$, $p = 0.14$; Figure 18D) connectivity profiles. We found a significant positive correlation between PVALB concentration and the negative BSR connectivity profile ($r = 0.273$, $p = 0.0028$, Bonferroni corrected $p = 0.0168$; Figure 18E) but not for positive profile ($r = -0.136$, p 0.14; Figure 18B). We did not see a significant relationship between HTR3A and the positive connectivity profile ($r = 0.189$, $p = 0.05$; Figure 18C). We did however observe a negative trend (i.e. not significant when corrected for multiple comparisons) for HTR3a expression and the negative connectivity profile ($r = -0.199$, $p = 0.0305$, Bonferroni corrected $p = 0.183$; Figure 18F).

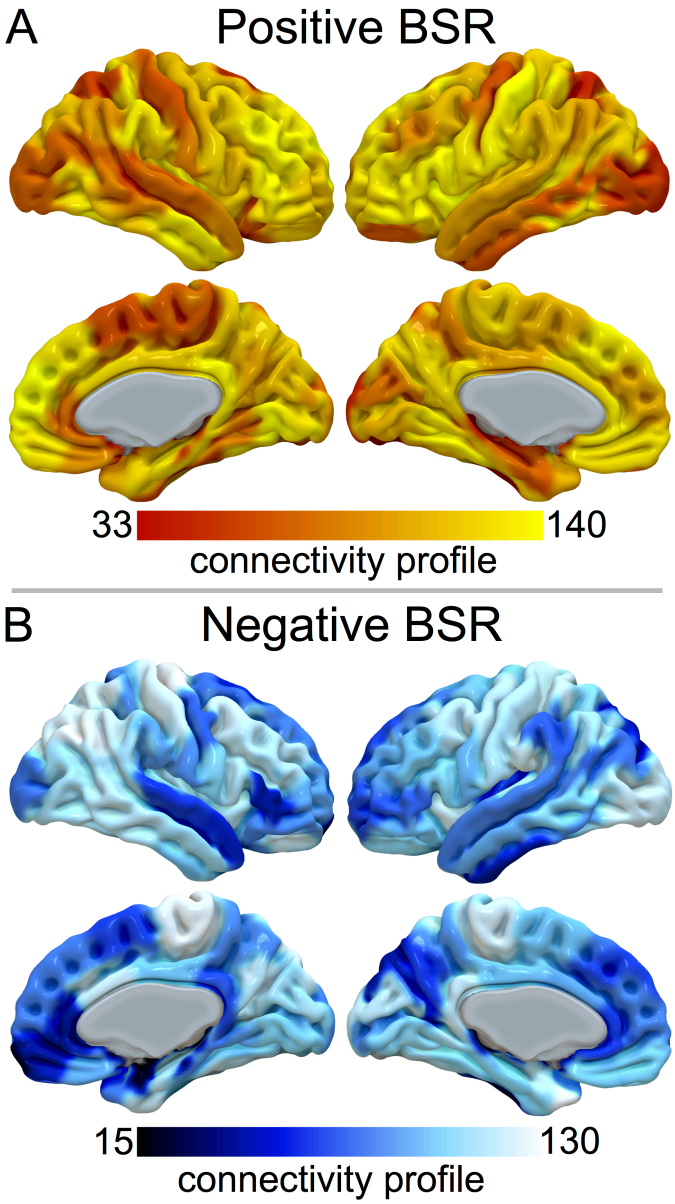

Figure 17. Regional connectivity for BSR scores during propofol administration
Each regional connectivity map was created by adding each region's BSR scores for each connection. Higher values correspond to that region having a high number and highly weighted BSR scores that either gain or lose connectivity during sedation. **A** shows regions that gain connections during sedation. **B** shows regions that lose connections during sedation.

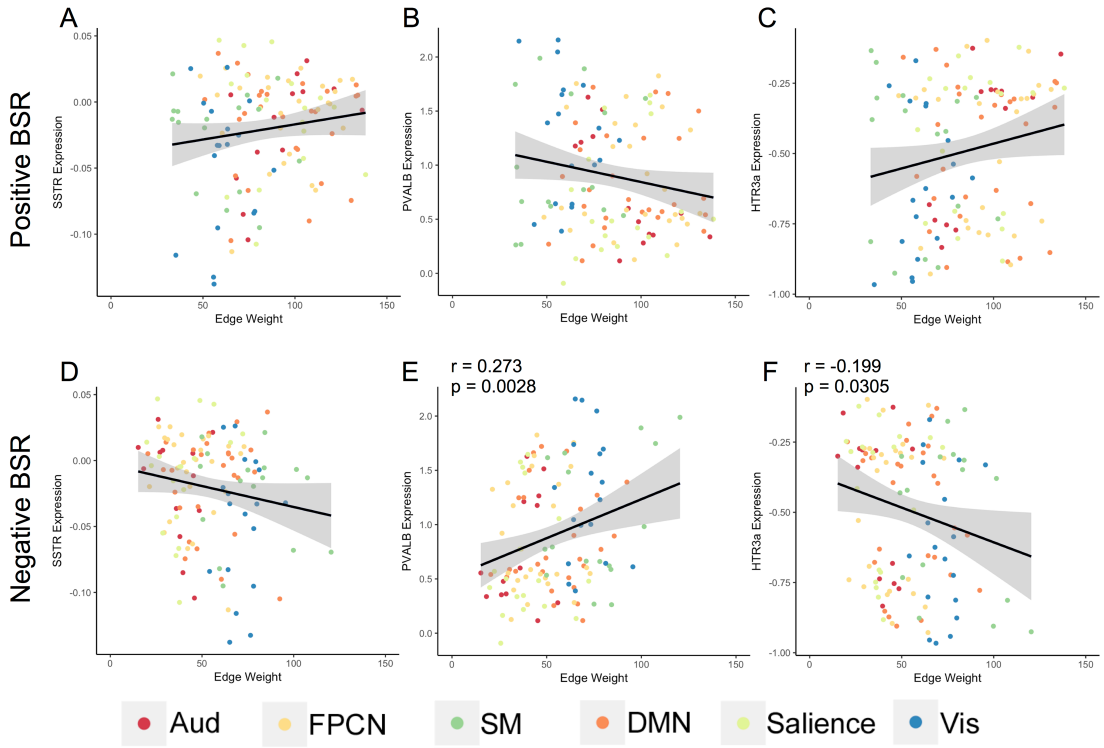

Figure 18. Correlating regional connectivity values with interneuron gene expression
These graphs show correlations with the regional connectivity values shown in Figure 17 with the cortical distribution of interneuron gene expression shown in Figure 13. The top row shows correlations with regions that increase connectivity during sedation and the bottom row shows correlations with regions that decrease connectivity during sedation. The first column shows correlations with somatostatin, the middle column shows correlations with parvalbumin and the third column shows correlations with 5-HT3a. **E** shows a significant positive correlation between regions that lose connectivity during sedation and parvalbumin expression ($r = 0.273$, $p = 0.0028$, Bonferroni corrected $p = 0.0168$). **F** shows a marginally significant (does not survive correction or multiple comparisons) negative correlation between regions that lose connectivity and expression of 5-HT3a ($r = -0.199$, $p = 0.0305$, Bonferroni corrected $p = 0.183$).

3.4. Discussion

The first aim of the present study was to identify patterns of cortical connectivity that are associated with propofol-induced anaesthesia. Using fMRI data from 25 healthy participants and the data-driven statistical technique PLS, we found a pattern of connectivity that corresponded to connections that either increased or decreased in connectivity strength during sedation induction, then returned back to their initial levels during recovery. We then investigated which large-scale cortical networks were most involved in sedation and found that the DMN displayed a highly dynamic connectivity profile with subnetworks that variably increased or decreased in

connectivity strength. We also found that internetwork connections between the Vis and SM networks significantly disconnected at higher levels of sedation, whereas internetwork connections between DMN and attentional and executive function networks, including the FPCN and SNs, increased during sedation. In addition to this connectivity analysis, we examined the distribution of the genes coding for PV, Sst and 5HT3a receptors on GABAergic interneurons. We found that the regions that disproportionately reduced in connectivity during sedation significantly overlapped with the pattern of PV expression in the cortex. Taken together, our findings suggest that propofol sedation significantly alters brain network connectivity patterns and that, in the cortex, this process is driven by PV expressing GABAergic interneurons.

We found that cortical network activity follows two dynamic patterns of connectivity, both of which are likely to disrupt information processing in the brain and ultimately lead to a loss of consciousness. Previous work has found that anaesthetic drugs cause sedation and unconsciousness by altering cortical integration (Alkire et al., 2008; Mashour, 2004). The present results are consistent with this hypothesis by finding that loss of integration is not simply the loss of connectivity between regions, but also the increase in connectivity with regions and networks that are anticorrelated in the awake, healthy human brain. This increase in connectivity between regions that are normally anticorrelated has been found both in anaesthesia administration and patients with disorders of consciousness. Stamatakis et al. (2010) used seed-to-voxel functional connectivity to measure changes in PCC connectivity, a primary node in the DMN, during propofol sedation. They found that PCC increased in connectivity with regions of ECN, including the anterior cingulate cortex (ACC) at deeper levels of sedation. They also found increased connectivity with precentral gyrus and pontine tegmentum. This effect has also been observed in monkeys undergoing anaesthesia (Barttfeld et al., 2015). Attention and executive function networks have also been found to be important in conscious information processing. Naci et al. (2013) found that activity changes in an attentional network reliably predicted whether healthy participants responded “yes” or “no” to a question. Together, this body of work suggests that DMN and attentional networks both play important roles in conscious processing.

Furthermore, increases in connectivity between networks that are normally anticorrelated have also been observed in patients with DOC. As previously described, Di Perri et al. (2016) examined functional connectivity in patients with UWS, MCS, eMCS and healthy control subjects and found that patients with more severe pathology had significantly reduced within DMN connectivity and greater between network connectivity with DMN and putative task positive regions. Our results reinforce these previous findings, suggesting that between DMN-TPN connectivity generally reflects a subject's level of consciousness and is not specific to a particular drug or pathology.

Our results show that within DMN connectivity is not simply decreased during propofol sedation, but that the DMN displays a dynamic, multifaceted pattern of connectivity with certain connections increasing and decreasing with sedation. Previous work has suggested that the low frequency functional architecture of the brain under anaesthesia, as measured by fMRI, remains relatively invariant (Vincent et al., 2007). Though this intrinsic organisation of brain networks is likely accurate at low frequencies, our inclusion of higher frequencies and the use of a data-driven method reveals that alterations in the dynamics of large-scale networks are present during sedation.

This is in line with previous work that has shown DMN regions, particularly posterior regions like PCC and precuneus, decrease in activity in response to propofol and sevoflurane administration (Kaisti, Oikonen, Aalto, Hinkka, & Scheinin, 2002). These regions have also been shown to play an important role in patients with disorders of consciousness, with evidence suggesting DMN regions typically disconnect in vegetative state patients, but are the first to reactivate in patients who recover (Laureys, Boly, & Maquet, 2006). Here we report that a significant number of connections within DMN follow one of two patterns. One pattern shows these connections increasing with sedation, the other shows them decreasing. It is possible that decreases in overall activity of primary DMN nodes do not necessarily result in reduced connectivity between all DMN regions. It may be the case that different DMN subsystems (Andrews-Hanna et al., 2010) follow different dynamics when primary nodes reduce in activation strength.

The network effects reported here are not restricted to higher-level cognitive networks, but also show that primary sensory networks alter their connectivity pattern during sedation. We found that the Vis and between Vis-SM connectivity reduced significantly during sedation, then strengthened during recovery. We also found that Auditory network connectivity increased with sedation, then decreased during recovery. This is consistent with previous work showing that Auditory network connectivity was most predictive of whether a patient was in a minimally conscious state or vegetative/unresponsive wakefulness syndrome (Demertzi et al., 2015). In contrast, other work shows that primary sensory cortices do not significantly change their connectivity pattern during propofol anaesthesia (Boveroux, et al. 2010; Schroter et al., 2012). Future work could examine connectivity patterns in primary sensory/motor cortices directly to more clearly identify whether they play a causal role is sedation or are indirectly affected by altered cortical and subcortical connectivity.

In addition to identifying patterns of connectivity associated with propofol-induced sedation, we were also interested in seeing whether different GABAergic interneuron subtypes were disproportionately expressed in regions greatly affected by propofol. This was motivated by previous work showing that propofol binds with high affinity to the β subunit of the GABAa receptor (Yip et al., 2013). Using the AIBS Human Microarray database we constructed expression maps for PVALB, SSTR and 5HT3a, receptors that mark the three subtypes of GABAergic interneurons found in the cortex (Rudy et al., 2013; Tremblay et al., 2016). We found that regions that lost connectivity correlated significantly with regions that contained a high concentration of PVALB expression. Though we cannot examine it directly due to the low sampling frequency of fMRI, this finding is potentially reflective of the loss of long-distance gamma coherence during sedation (Alkire et al., 2008; John, Prichep, Aubert, & Gugino, 2001; Lewis et al., 2012). Gamma oscillations are thought to be essential for binding neuronal assemblies and playing an important role in various cognitive processes including memory, sensation and influencing emotional states (Carr, Karlsson, & Frank, 2012; Meletis et al., 2009; Phelps, 2006). It is well established that interconnected parvalbumin interneurons induce gamma oscillations through multiple means (Buzsaki, 1983; Buzsaki & Wang, 2012; Sohal et al., 2009). Thus, our finding that regions that lose long distance connectivity overlap with regions with

high concentrations of PVALB expression may be reflective of a local circuit mechanism where propofol binds to GABAergic receptors normally mediated by parvalbumin interneuron activity. Recent work by Huang et al. (2016) found that decreasing inhibitory input to parvalbumin interneurons, leading to increased but desynchronised firing, results in desynchronised firing of pyramidal cells and ultimately an impairment of gamma oscillations. Propofol may induce a similar effect, where it stochastically binds to GABAergic receptors on postsynaptic pyramidal neurons, resulting in a loss of synchrony and therefore a loss in long-range gamma oscillations.

Chapter 4. Structure-function relationships between default mode and salience networks during propofol-induced sedation

4.1. Introduction

As demonstrated in the previous chapter, anaesthesia-induced alterations in conscious experience have been shown to be associated with significant alteration in brain network connectivity. These changes have been observed using both functional connectivity measured with fMRI (Bonhomme et al., 2016; Boveroux et al., 2010; Stamatakis et al., 2010) and EEG (Chennu et al., 2016). As demonstrated in the previous chapter and elsewhere, network alterations induced by propofol have been shown to result in the breakdown of several neural systems (Hudetz, 2012). Specific large-scale brain networks, including the DMN (Stamatakis et al., 2010) and FPCN (Lee et al., 2013), have also been shown to alter their connectivity during propofol-induced anaesthesia. The DMN has been shown to play a significant role in internally directed attention (Andrews-Hanna et al., 2010) and in mediating certain elements of task performance (Vatansever et al., 2017). The FPCN plays a role in cognitive control and executive function tasks (Dosenbach et al., 2006) and in mediating the interactions between itself and DMN and SN (Dixon et al., 2017; Spreng et al., 2013). The interactions between these three higher-level cognitive brain networks have been shown to be transiently anticorrelated at low frequencies (Fox et al., 2005; Dixon et al. 2017) and altered during sedation (Boveroux et al., 2010). However, the stability of these relationships has not been directly examined under propofol sedation. The previous chapter showed both increased and decreased connectivity between regions within these networks during sedation. However, recall that the bootstrap ratio is only sensitive to relative changes in r-values and does not provide information on whether specific connections switch from being anticorrelated to positively correlated between conditions.

A deeper characterisation of connectivity under anaesthesia would help identify whether certain networks are more associated with on-going cognitive

processing or the brain's intrinsic functional architecture. In this vein, Barttfeld et al. (2015) used a non-human primate model based on an experiment in which macaques underwent propofol-induced anaesthesia to further understand these connectivity patterns and examine brain network structure-function relationships during sedation. The results showed that during anaesthesia, whole-brain functional connectivity lowered information capacity, reduced network anticorrelations, and increased overlap with structural connectivity. These findings suggest that the functional connectivity observed during wakefulness reduces in variability and regresses back to a pattern dictated by the underlying structural framework during sedation. These results further suggest that a signature of consciousness is related to how complex the pattern of connectivity is and how it is related to structural connectivity. Though this study has shaped our current understanding of functional connectivity during sedation, it is limited by its use of an animal model and its focus on whole-brain connectivity. No study has yet examined the relationship amongst direct structural connections between inter and intra-network regions and how they mediate functional connectivity patterns at different levels of consciousness in humans.

Therefore, the study reported in this chapter aimed to understand the relationship between the DMN, the FPCN and the SN during propofol sedation as well as examine the role of structural connections in mediating their connectivity. First, we measured the strength of functional connectivity both within and between each of these three large-scale brain networks. Next, we used graph theoretical measures to model various topological characteristics of within and between network functional connectivity. Finally, using diffusion-weighted structural imaging, we measured the degree of coherence between structural and functional connectivity.

We hypothesised that resting state connectivity between the DMN and the putative TPNs prior to propofol anaesthesia would be very low or anticorrelated. As previous work in patients with DOC has shown that anticorrelation between networks becomes positive in unconscious individuals (Di Perri et al., 2016), we also hypothesise that this pattern of increased functional connectivity would be found while participants were undergoing propofol-induced sedation. Furthermore, if the overlap between structure and function between large-scale networks also increased

during sedation, it would suggest that functional connections between networks become less flexible and increasingly rely on those underlying white matter tracts.

4.2. Methods

4.2.1. Participants

For information on the participants in this study, refer to section 3.2.1 in Chapter 3.

4.2.2. Sedation Procedure

For details on the sedation procedure, refer to section 3.2.2 in Chapter 3.

4.2.3. Data Acquisition

4.2.3.1. T1-weighted and fMRI Data

For details on the acquisition of T1-weighted and fMRI data, refer to section 3.2.3 in Chapter 3.

4.2.3.2. Diffusion Tensor Imaging Data

To examine the relationship between structure and function in large-scale brain network connectivity, we constructed a standardised connectome using High Angular Resolution Diffusion Imaging (HARDI) from the Human Connectome Project following methods defined in Van den Heuvel and Sporns (2011). Data was collected from 50 healthy young adult participants all between the ages of 20 and 35. Using HARDI data is advantageous for identifying complex axonal morphologies due to the fact the diffusion-weighted signal is acquired in many directions, thus providing increased accuracy in determining precise axonal morphology (Setsompop et al., 2013). Data were acquired using a single shot single refocusing spin-echo, echo-planar imaging sequence with a multishell diffusion scheme of 90 diffusion-weighted directions with b-values of 990, 1985, and 2985 s/mm^2 and spatial resolution of 1.25mm^3.

4.2.4. Data Pre-processing

4.2.4.1. Brain Network Parcellation Scheme

To construct brain networks, we used the 129 region Lausanne Parcellation (Hagmann et al., 2008), which consists of 64 ROIs per hemisphere and 1 bilateral brain stem mask. In order to assign each ROI to one of the networks (after making sure that all the images are in the same space), we calculated the overlap for each ROI with each sub-network. The maximum overlap was used to assign the ROI to the respective sub-network. As this chapter focuses on high-level cognitive networks, we only included the regions from the DMN (27 regions), FPCN (26 regions) and SN (18 regions).

4.2.4.2. High-resolution T1-weighted and Resting State Image Pre-processing

For details regarding pre-processing, see 3.2.4 for T1-weighted and resting state images, refer to section 3.2.4 in Chapter 3.

4.2.4.3. Diffusion Data Pre-processing

The diffusion data were processed using generalized q-sampling imaging (GQI) in DSI Studio (http://dsi-studio.labsolver.org/). GQI quantifies the density of diffusion in different orientations, providing directional information regarding crossing fibres. Its use of the spin distribution function (SDF) provides greater sensitivity and specificity than previous methods at identifying white matter tracts (Yeh, Wedeen, & Tseng, 2010). We used a standard diffusion-sampling threshold of 1.25 with a deterministic fibre-tracking algorithm (Yeh, Verstynen, Wang, Fernández-Miranda, & Isaac, 2013). The anisotropy ratio was 0.0518329, the angular threshold was 60 degrees, and the step size was 0:625 mm[3]. Tracks with a length of less than 10mm were discarded. A 71 x 71 connectivity matrix was created by using each subject's Lausanne parcellation and quantifying Generalized Fractional Anisotropy (GFA) between each brain region. GFA uses spherical harmonics and Orientation Distribution Functions (ODF) to more accurately characterise estimates of anisotropy than earlier methods (Cohen-Adad et al., 2008; Tuch, 2004). The group average connectivity matrix was

[3] These are the default settings for GQI implemented in DSI Studio.

constructed by selecting connections (as defined by a positive gfa value) that were present in at least 75% of the HCP subjects (van den Heuvel & Sporns, 2011). This matrix was then binarised and used as a healthy group averaged connectome for comparison to the Sedation functional data.

4.2.5. Functional Connectivity

Functional connectivity was calculated by measuring the correlation coefficient between each of the 71 regions in the DMN, the FPCN and the SN networks. This allowed us to examine both within- and between-network connectivity for both the Sedation and DOC datasets. These measures were calculated by averaging the correlation values within or between each network of interest. We did not threshold these network measures since our hypotheses stipulate the importance of anticorrelations between these networks.

4.2.6. Function-to-Structure Overlap

We modified a measure of function to structure overlap initially defined by Barttfeld et al. (2015). We calculated the similarity of functional and structural matrices by measuring the correlation coefficient between the vectorised structural matrix and the vectorised functional matrix. The difference between Barttfeld et al. (2015) and our approach is that we used the raw functional connectivity matrix, whereas they used the output from a clustering analysis. We made this change because Barttfeld et al. (2015) did not order their nodes based on pre-defined brain networks, and therefore they required a way to standardise the node order across each matrix. As we used a pre-defined set of brain regions, our nodes were ordered prior to analysis.

4.2.7. Graph Theoretical Metrics

4.2.7.1. Graph Construction

Brain graphs were constructed using the python library Brainx (https://github.com/nipy/brainx) based on the Networkx graph theory package (https://networkx.github.io/). To generate the sparse graphs necessary for these analyses, we calculated each metric over a range of thresholds (top 10%, 15%, 20%,

and 25% of connections kept) and averaged them to get the final metric (Bullmore & Bassett, 2011).

4.2.7.2. Network Modularity and System Segregation

We calculated graph modularity and system segregation to assess how compartmentalised processing was amongst the DMN, FPCN and SN. Modularity (*Q)* is defined as the number of within-network connections compared to all connections and quantifies the strength of segregation of distinct networks. It is defined as:

$$Q = \sum_{i=1}^{m} (e_{ii} - a_i^2)$$

Where e_{ii} is the fraction of edges that connect nodes within module *i,* a_i is the fraction of edges that connect a node in module *i* to any other node, and *m* is the number of modules.

System segregation is a measure of network segregation based on nodal assignment to a specific network (Chan, Park, Savalia, Petersen, & Wig, 2014). It is defined as the mean connectivity strength of edges between each pair of nodes within the same network ($\bar{z}_w$*).* Between-network strength is calculated as the mean connectivity between nodes of two different networks ($\bar{z}_b$). The system segregation metric is calculated by subtracting mean between-network connectivity from mean within-network connectivity and dividing by mean within-network connectivity:

$$System\ Segregation = \frac{\bar{z}_w - \bar{z}_b}{\bar{z}_w}$$

4.2.7.3. Network Efficiency

Efficiency is calculated by using the minimum path length (*L*), which is the smallest number of connections that must be crossed to get from node *i* to node *j*. Nodal efficiency is calculated by taking the inverse of the harmonic mean of *L* between node *i* and all other nodes:

$$E_{nodal}(i) = \frac{1}{(N-1)} \sum_{j \in G} \frac{1}{L_{i,j}}$$

Local efficiency, a measure of network segregation, is a measure of information transfer efficiency within neighbouring nodes (i.e. nodes with direct edges). It is calculated as the average nodal efficiency in nodes that directly connect to node *i*, excluding node *i* itself:

$$E_{local} = \frac{1}{N_{G_i}(N_{G_i} - 1)} \sum_{j,k \in G_i} \frac{1}{L_{j,k}}$$

Global efficiency measures the efficiency of information transfer among all pairs of nodes in a graph. It is calculated as the average nodal efficiency of all nodes, as follows:

$$E_{global} = \frac{1}{N(N-1)} \sum_{i \neq j \in G} \frac{1}{L_{i,j}}$$

4.2.7.4. Network Hub Measures

We also assessed whether each node in the network worked as a hub for routing information between or within large-scale networks. To do so, we calculated the participation coefficient (PC) and the within-module degree (WD) for each node. The PC is a measure of internetwork connections and is defined as:

$$PC_i = 1 - \sum_{s=1}^{N_M} \left(\frac{k_{is}}{k_s}\right)^2$$

Where k_{is} is the degree of node *i* to other nodes within its own network (*s*), and k_i is the degree of node *i* to any node regardless of network assignment. Subtracting that ratio from 1 results in a normalised measure.

WD can be defined as:

$$WD_i = \frac{k_{is} - \bar{k}_{s_i}}{\sigma_{k_{s_i}}}$$

Where k_{is} is the degree of node *i* to nodes in its own network (*s*), $\bar{k}_{s_i}$ is the average degree of all nodes in the network *s*, and $\sigma_{k_{s_i}}$ is the standard deviation of the degree of

all nodes in *s*. WD is a standardised metric for how connected a node is to other within-network nodes.

These two measures can then be used to characterise each node in the network as either a connector or provincial hub. The number of connector hubs is defined as the number of nodes important for between-network, and not within-network communication (Guimera & Nunes Amaral, 2005). Nodes that surpass WD < 2.5 and PC > 0.62 are classified as connector hubs. Provincial hubs are important for within network, but not between network communication. They are defined as nodes with WD > 2.5 and PC < 0.3.

4.2.8. Statistical Analysis

A paired t-test was used to test for significant differences between awake and moderate sedation conditions in the 25 healthy subjects in the Sedation dataset. As we were interested in finding gross differences between normal wakefulness and sedation, we chose to only include the awake and moderate sedation conditions in this analysis. All significant p-values reported were Bonferroni corrected.

4.3. Results

4.3.1. DMN, FPCN and SN Functional Connectivity in Sedation

For the sedation condition, we did not find significant differences in within network connectivity for DMN ($t_{(24)} = 1.299$, $p = 0.206$; Figure 19A) or FPCN ($t_{(24)} = -1.67$, $p = 0.107$; Figure 19B). Functional connectivity in the SN was greater in moderate sedation compared to awake conditions, however this effect did not survive multiple comparisons ($t_{(24)} = -2.08$, $p = 0.048$, Bonferonni corrected $p = 0.77$; Figure 19C).

For between network connectivity in the sedation condition, we did not find any significant differences in strength for between DMN-FPCN ($t_{(24)} = -1.39$, $p = 0.177$; Figure 20A), and FPCN-SN connectivity ($t_{(24)} = -2.005$, $p = 0.056$; Figure 20C). However, we did detect significantly greater connectivity during the sedation condition compared to awake condition between DMN and SNs ($t_{24)} = -3.371$, $p = 0.0025$, Bonferonni corrected $p = 0.0404$; Figure 20B). Further, we averaged

connectivity strength across participants within subjects and found that the number of anticorrelations decreased substantially in the sedation condition (17 anticorrelated edges) compared to the awake condition (68 anticorrelated edges) (Figure 20D and E; graphs thresholded at the whole brain level with top 20% positive correlations and top 20% negative correlations).

4.3.2. Function-Structure Overlap in Sedation

For the sedation condition, we found no significant differences in function-structure overlap within DMN ($t_{(24)}$ = 0.23, p = 0.812; Figure 21A), within FPCN ($t_{(24)}$ = 0.0145, p = 0.988; Figure 21B), or within SN ($t_{(24)}$ =0.954, p = 0.349; Figure 21C). For between network function-structure overlap, we did not see any significant differences in the sedation condition between DMN-FPCN ($t_{(24)}$ = -0.978, p = 0.337; Figure 22A) or between FPCN-SN ($t_{(24)}$ = -2.005, p = 0.056; Figure 22C). However, we did observe significantly greater overlap between DMN-SN in the sedation condition in comparison to the awake condition ($t_{24)}$ = -3.474, p = 0.0019, Bonferonni corrected p = 0.031; Figure 22B). We have also visualized overlap of the underlying structural connectivity between DMN and SN in awake (Figure 22D) and sedation (Figure 22E) conditions, with blue edges corresponding to white matter connections defined using HARDI DTI and pink connections corresponding to a significant functional connection overlapping with a structural connections.

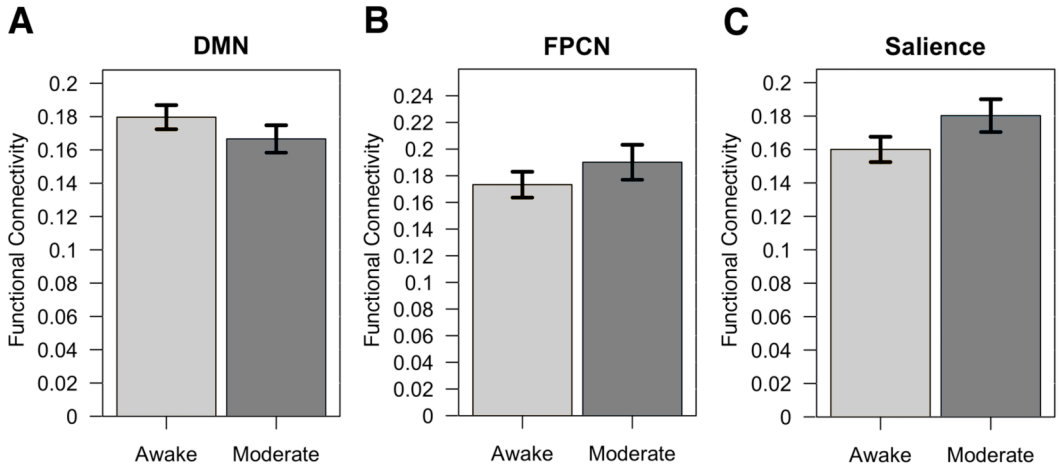

Figure 19. Functional connectivity within networks during sedation
A shows within DMN connectivity ($t_{(24)}$ = 1.299, p = 0.206), **B** shows within FPCN connectivity ($t_{(24)}$ = -1.67, p = 0.107), and **C** shows within SN connectivity ($t_{(24)}$ = -2.08, p = 0.048, Bonferonni corrected p = 0.77) between awake and moderate sedation conditions. Error bars reflect standard error.

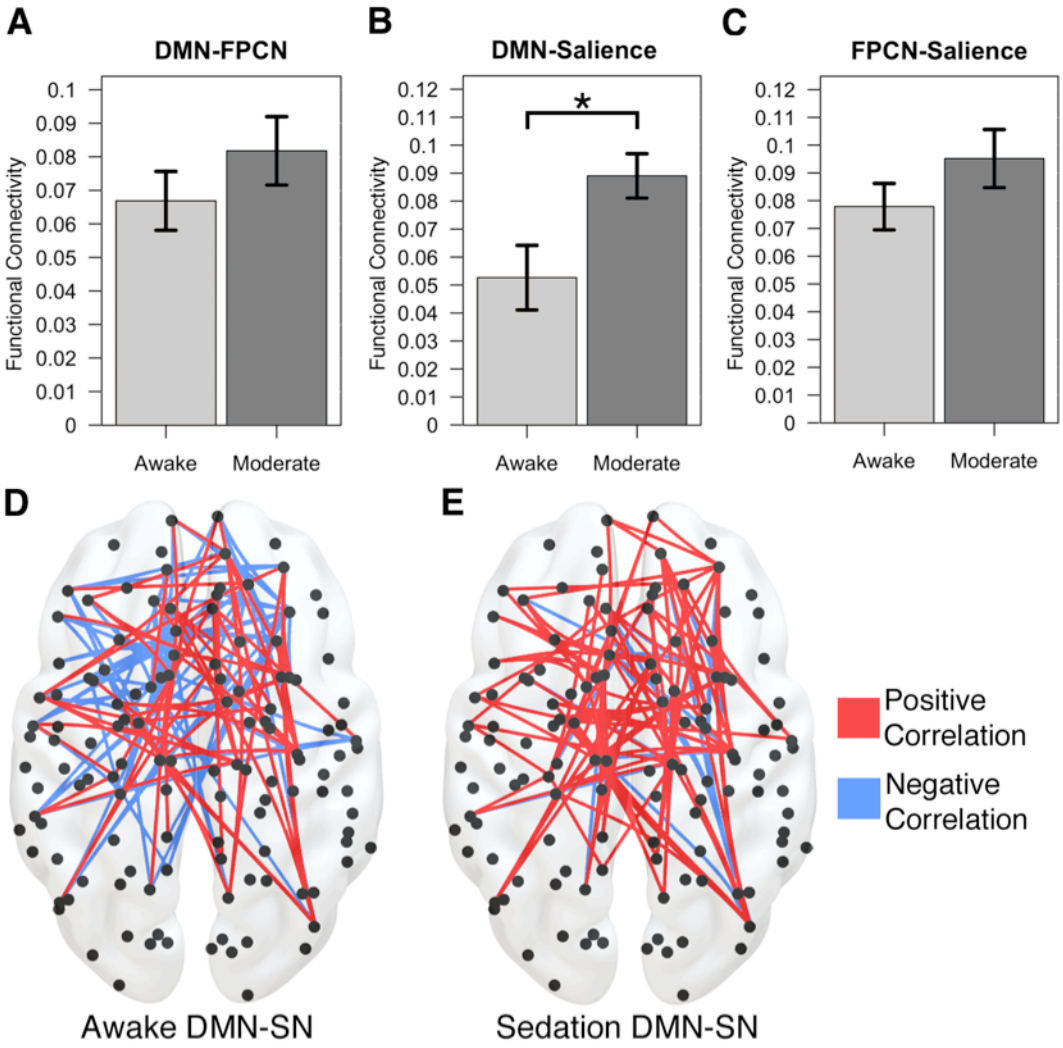

Figure 20. Functional connectivity between networks during sedation

A shows between DMN-FPCN connectivity ($t_{(24)}$ = -1.39, p = 0.177), **B** shows a significant difference in between DMN-SN connectivity ($t_{24)}$ = -3.371, p = 0.0025, Bonferonni corrected p = 0.0404) and **C** shows between FPCN-SN connectivity ($t_{(24)}$ = -2.005, p = 0.056) between awake and moderate conditions. Error bars reflect standard error. **D** Positive and negative correlations between DMN and SN connections during the awake condition (thresholded at top 20% positive correlations and top 20% negative correlations). **E** Positive and negative correlations between DMN and SN connections during the sedation condition (thresholded at the whole brain level with top 20% positive correlations and top 20% negative correlations).

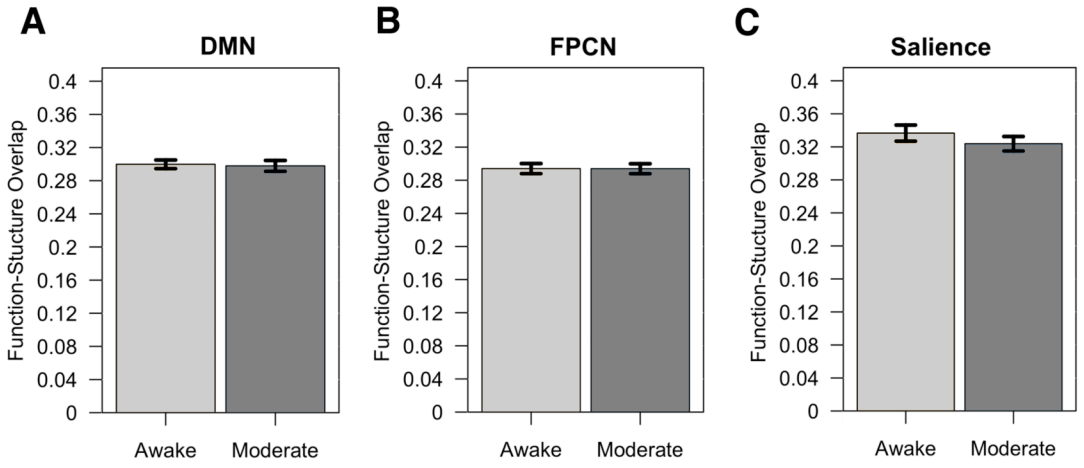

Figure 21. Function to structure overlap within networks during sedation
A shows within DMN function-structure overlap ($t_{(24)}$ = 0.23, p = 0.812), **B** shows within FPCN function-structure overlap ($t_{(24)}$ = 0.0145, p = 0.988), and **C** shows within SN function-structure overlap ($t_{(24)}$ =0.954, p = 0.349) between awake and moderate sedation conditions. Error bars reflect standard error.

4.3.3. Integration and segregation graph theoretical measures in sedation

As formal measurements of integration and segregation are thought to be important in conscious processing (Tononi et al., 2016), we used a set of graph theoretical measures specifically aimed at assessing these (Cohen & D'Esposito, 2016) to determine whether they could differentiate between wakeful and sedation conditions.

We did not observe any differences in either global efficiency ($t_{(24)}$ = 2.309, p = 0.0298, Bonferroni corrected p = 0.477; Figure 23A) or the number of connector hubs ($t_{(24)}$ = -0.188, p = 0.853; Figure 23B), two measures of network integration. We did observe significantly greater system segregation in the awake condition compared to the sedation condition ($t_{(24)}$ = 3.524, p = 0.0017, Bonferroni corrected p = 0.0277; 23A). We did not see any differences in modularity ($t_{(24)}$ = 1.694, p = 0.103; Figure 23B), local efficiency ($t_{(24)}$ = -1.512, p = 0.147; Figure 23C), or the number of provincial hubs ($t_{(24)}$ = -1.112, p = 0.276; Figure 23D).

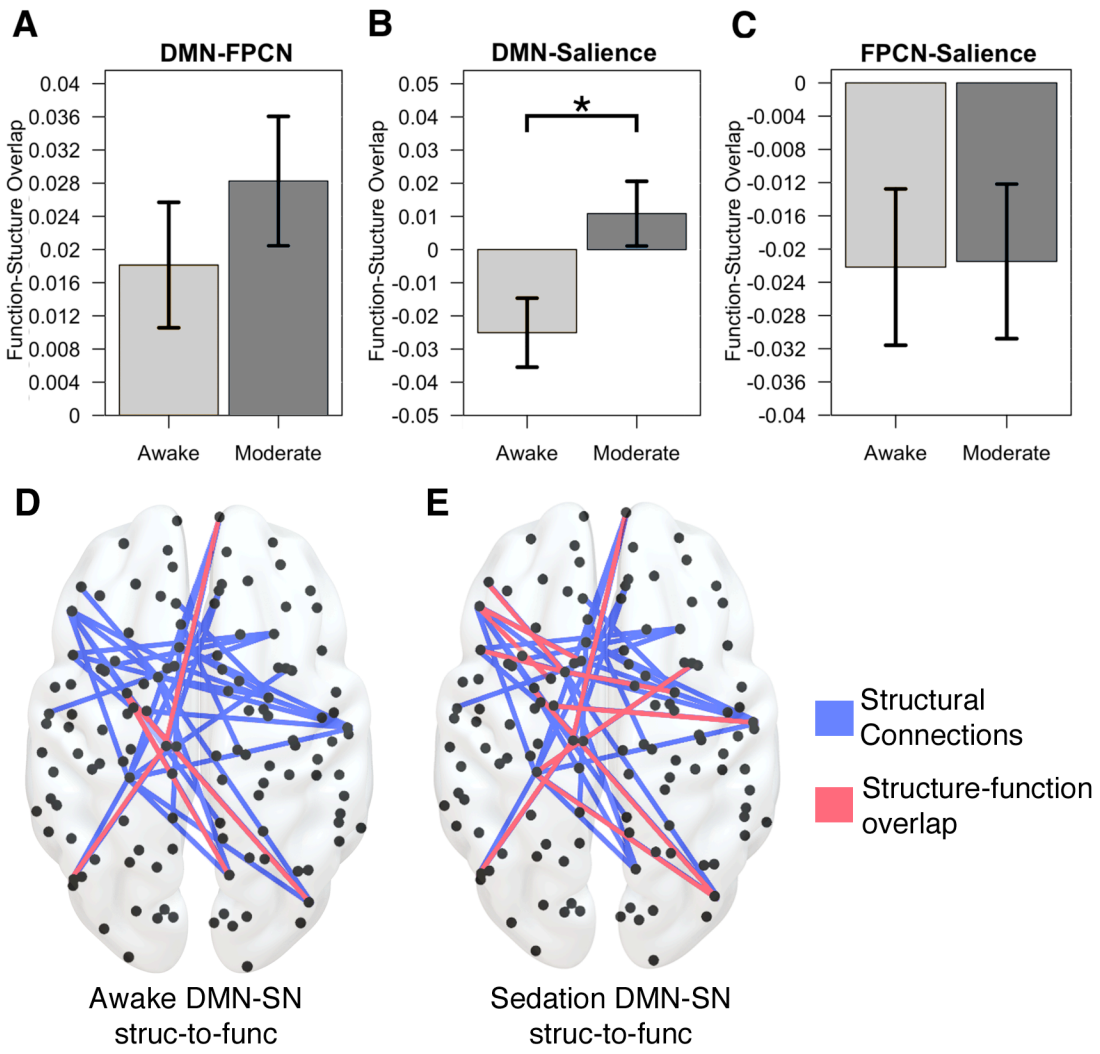

Figure 22. Function to structure overlap between networks during sedation
A shows between DMN-FPCN function-structure overlap ($t_{(24)}$ = -0.978, p = 0.337), **B** shows a significant difference in between DMN-Salience function-structure overlap (t_{24} = -3.474, p = 0.0019, Bonferonni corrected p = 0.031) and **C** shows between FPCN-Salience function-structure overlap ($t_{(24)}$ = -2.005, p = 0.056) between awake and moderate conditions. Error bars reflect standard error. **D** shows non-overlapping structural connections (blue) and overlapping structural connections (pink) with functional connections during the awake condition. **E** shows non-overlapping structural connections (blue) and overlapping structural connections (pink) with functional connections during the sedation condition.

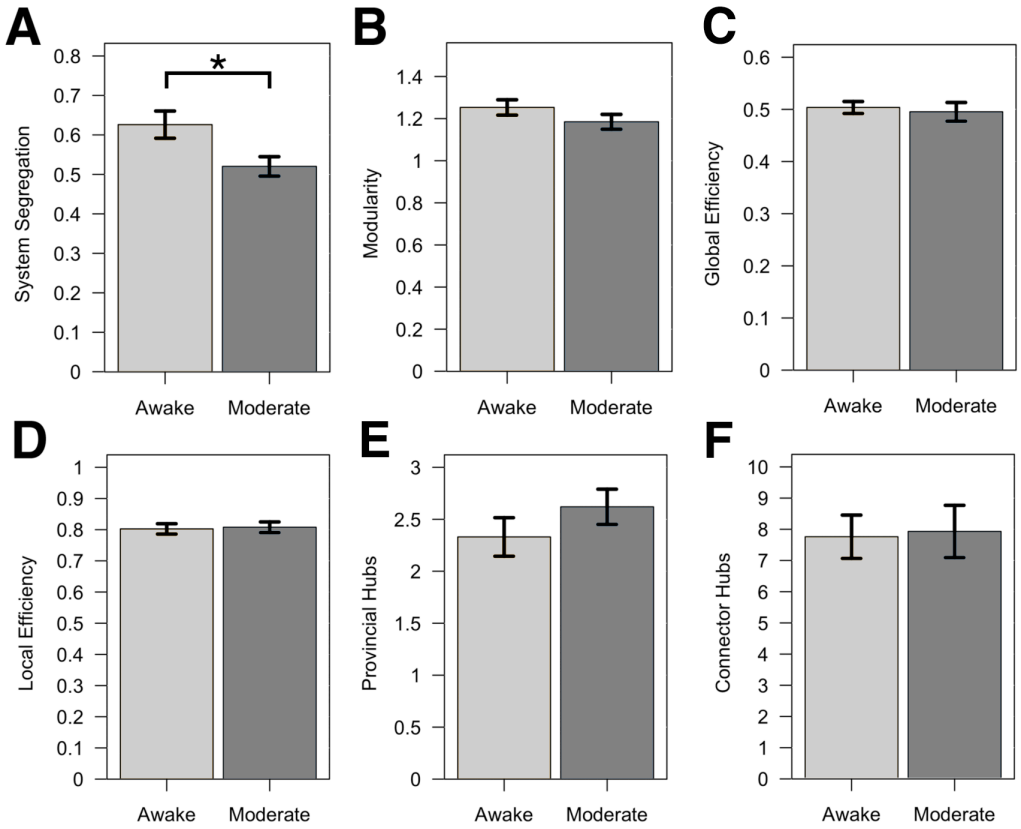

Figure 23. Graph theoretical measures between awake and moderate sedation conditions
A shows a significant difference in system segregation ($t_{(24)} = 3.524$, p = 0.0017, Bonferroni corrected p = 0.0277) between awake and moderate sedation conditions. **B** shows modularity ($t_{(24)} = 1.694$, p = 0.103), **C** shows global efficiency ($t_{(24)} = 2.309$, p = 0.0298, Bonferroni corrected p = 0.477). **D** shows local efficiency ($t_{(24)} = -1.512$, p = 0.147) and **E** shows the number of provincial hubs ($t_{(24)} = -1.112$, p = 0.276) between awake and moderate sedation conditions. **F** shows the number of connector hubs ($t_{(24)} = -0.188$, p = 0.853)

4.4. Discussion

The present study focused on understanding the interactions within and between DMN, FPCN, and SN, motivated by the results obtained in Chapter 3, and their relationship to underlying structural connectivity during propofol-induced sedation. We found that during propofol administration, functional connectivity between DMN and SN increased significantly. Furthermore, the overlap between structural and functional connectivity increased between these two networks, suggesting a decreased repertoire of connectivity and as such an increased reliance on direct white matter tracts during sedation. Finally, we found that system segregation, a graph theoretical measure that characterises the degree of partitioning between brain networks, was significantly greater during wakefulness. In other words, the increased repertoire of functional connectivity patterns facilitates optimal segregation and supports wakefulness. This is consistent with previous work showing that sedation results in

decreased complexity and flexibility of brain network dynamics (Vincent et al., 2007; Barttfeld et al., 2015).

Our results showing an increase in connectivity (and conversely a decrease in anticorrelation) between DMN and SN expand upon previous work showing an increase in connectivity between DMN and putative task-positive brain regions. Stamatakis et al. (2010) showed an increase in connectivity between PCC, a primary node in the DMN, and several regions associated with attentional processes including the anterior cingulate cortex, an important region in the SN. Our results are consistent with this finding and expand it by demonstrating that this effect is not restricted to these two brain regions specifically, but can be observed between entire networks. Our finding also clarifies previous work using seed-to-voxel functional connectivity from DMN nodes in patients with DOC. Di Perri et al. (2016) found that patients who were in deeper states of clinical unconsciousness had reduced DMN anticorrelations compared to patients in more wakeful states and healthy controls. By demonstrating that this effect occurs specifically between DMN and SN, our study lays the groundwork for future research in DOC patients aimed at characterising the interaction between these two networks.

In addition to observing increased functional connectivity between DMN and SN during sedation, we found that the overlap between functional connectivity patterns and the structural connectivity between these networks increased. This finding confirms and expands on previous work in non-human primates (Barttfeld et al., 2015) by showing an increase in structure-function overlap at the whole-brain level in healthy human participants. The work also complements a recent study by Demertzi et al. (2019) that show increased structure-function overlap in patients with DOC. It also expands upon this work by examining structure-function overlap within and between specific large scale brain networks. This is advantageous in that it provides increased understanding of which networks underlie the neurophysiological processes associated with loss of consciousness. Future work could perform a similar analysis in DOC patients to examine whether the DMN-SN structure-function relationship is associated with lower levels of consciousness.

Uddin et al. (2009) used resting state functional connectivity with granger causality to show to that the PCC and vmPFC subregions of the DMN anticorrelate with different attentional and motor systems in awake healthy participants. Intriguingly, a Granger causality analysis showed that DMN regions exerted more influence on regions it is anticorrelated with than the other way around, suggesting that the DMN may modulate activity in attentional systems (Uddin, Kelly, Biswal, Castellanos, & Milham, 2009). Propofol administration may diminish the DMN's capacity to remain partitioned, resulting in increased activity traversing the structural connections between normally anticorrelated or weakly correlated networks.

We also found that system segregation, a measure of how distinct connectivity patterns are within and between networks, was greater during the awake state compared to the sedated state. These findings are also reflected in previous work modelling resting state functional connectivity with dynamical systems simulations (Bell & Sejnowski, 1995; Deco, Jursa, McIntosh, Sporns, & Kotter, 2009; Deco, Kringelbach, Jirsa, & Ritter, 2017). This work demonstrates that when coupling strength between regions is low, a state similar to sedation (Barttfeld et al., 2015; Jacobs et al., 2013) connectivity remains but is largely reflected by a single stable pattern that reflects the underlying structural connections. When coupling strength between regions is increased, the number of states increases, resulting in a multistable system. As the awake state is characterised by dynamic processing in multiple sensory, cognitive, and motor systems, it follows that it requires many stable, weakly coupled (or in some cases anticorrelated) networks. This stands in contrast to the sedated state, whereas unconsciousness approaches, less information is processed requiring fewer states and therefore less system segregation.

It is noteworthy that during the awake state, we found that average connectivity strength between DMN and SN was weakly positive. At first glance, this may seem to be in contention with previous work (Fox et al., 2005), however, as was demonstrated in Chapter 2, when examining between network interactions, it is important to highpass, and not bandpass filter functional data. Additionally, work by Dixon et al. (2017) used dynamic functional connectivity analysis to show that anticorrelations between these networks are transient during rest and can vary

significantly between individuals. Additionally, by visualizing these connectivity patterns, we found an almost equal number of positive correlations and anticorrelations during the awake state, further emphasizing the complexity of the connectivity patterns between these networks, and pushing this framework farther away from a simple adversarial relationship.

Chapter 5. Deep graph convolutional neural networks identify FPCN and DMN contributions to mental imagery

5.1. Introduction

The previous two chapters focused on changes in brain network connectivity while healthy participants underwent propofol sedation. We found connectivity within and between large-scale brain networks that underlie higher cognitive processes were significantly altered at higher levels of propofol-induced sedation. These findings echo previous work that examined resting state networks while participants underwent sedation (Boveroux et al., 2010; Stamatakis et al., 2010). They also reflect work examining brain networks in patients with DOC, which suggests that alterations within and between DMN are significantly altered in patients with more severe pathology (Di Perri et al., 2016; Vanhaudenhuyse et al., 2010). Since an aim of this thesis is to understand differences in drug-specific versus state-specific characteristics of consciousness, the final experimental chapter will apply several analyses from previous chapters to DOC patients. This chapter will focus on the development of an additional computational method in healthy participants to be used in the final chapter on DOC patients.

Due to the difficulty in diagnosing this patient group, machine learning-based computer-aided diagnosis of patients with DOC is an emerging and important area of research. A comprehensive review of machine learning is outside the scope of this thesis, however, some basic concepts are important to discuss to provide context for the choices of algorithms and metrics in the present chapter and Chapter 6. Briefly, machine learning is a field spanning mathematics and computer science which studies algorithms that can learn from data (Bishop, 2006). Broadly speaking, machine learning algorithms can be divided into those trained in a supervised or unsupervised manner. In supervised learning, the algorithm has access to data points that have associated labels or targets that it must learn. In unsupervised learning, the algorithm does not have targets or labels attached to its input data, and therefore must automatically find patterns in the structure of the data (Goodfellow, Bengio, &

Courville, 2016). All algorithms used in Chapters 5 and 6 are examples of supervised learning algorithms. Supervised learning can be further broken down into classification and regression models. Classification seeks to classify each data point as belonging to a set of predefined categories or bins, whereas regression attempts to predict the specific continuous variable[4] (Murphy, 2012). Chapters 5 and 6 will both use classification to categorise brain graphs as being in one of two possible conditions. Another important concept is data splitting. Datasets are typically split into training, validation and testing sets. The training set is the largest set and is fed into the algorithm during training, allowing it to learn by example and reduce its training error. The validation set is used to "tune" the algorithm, by allowing it to be tested on small amounts of data not contained in the training set. Finally, the test set is entirely held out from training and is used to assess how well the algorithm will perform on data that it has not yet seen (Goodfellow et al., 2016).

A critical challenge in machine learning is developing a model that does not just perform well on the training and validation sets, but that it can *generalise* to the test set. The degree to which the algorithm accurately predicts targets in the test set is known as the generalisation error. If a model does not perform well on any dataset, it is known as *underfitting*. When a model performs well on the training set but has a high generalisation error, it is said to be *over-fitting* the training set (Bishop, 2006). Both situations require the tuning of the algorithm's hyperparameters, which are settings that help control the algorithm's behaviour. Each algorithm has its own set of tunable hyperparameters. Different algorithms necessitate different optimal approaches to tuning, as it can be very time consuming to tune hyperparameters with continuous values by hand (Snoek et al., 2015).

There are many types of machine learning algorithms, each of which having their own set of advantages and disadvantages. Broadly, algorithms can be categorised as those that require features to be engineered prior to input and those that engineer their own features from raw data (also known as deep learning models) (Goodfellow

[4] A relevant example of a classification problem is categorizing a patient as having MCS or UWS. An example of a regression model would be attempting to predict a patients CRS-R score.

et al., 2016). Algorithms that require features to be engineered, like logistic regression or support vector classifiers, typically train much faster and require less data than deep learning models. A feature can be any measurable property of an observed phenomenon[5] (Bishop, 2006). Several methods, including resting state functional connectivity, have been used to produce input features to classifiers with the aim of diagnosing patients utilising brain connectivity information (Chennu et al., 2017; Demertzi et al., 2015). One possible limitation of this work is that it relies on a clinical diagnosis in order to label each patient's data as belonging to a specific category (e.g. unresponsive wakefulness syndrome or minimally conscious state). This means that the classifier will learn any error in the clinical diagnosis of these patients, which could result in the classifier simply reiterating biases that are already inherent in clinical diagnoses. With this in mind, Chapter 6 will stratify DOC patients not by clinical diagnosis, but by their capacity to perform a volitional mental imagery task while in the scanner (Monti et al., 2010; Owen et al., 2006). The classifiers will then be trained on the patient's resting state data with the goal of differentiating patients who have the ability to perform a volitional task from those who cannot.

Another possible limitation of previous work in developing classifiers for DOC patients is that they rely on more classical feature engineering-based machine learning methods such as support vector machines. Recent years have seen an explosion of research in deep learning with artificial neural networks (Lecun, Bengio, & Hinton, 2015), a subfield of machine learning that has proven to be highly successful in many research and industrial applications. In deep learning, raw data is fed into the algorithm, which then extracts features of the data that it has determined are useful for a given goal (i.e. classification or regression). It does this by stacking multiple layers of nonlinear processing units (known as neurons) that extract multiple levels of representation for each incoming data point (Goodfellow et al., 2016). One successful type of deep learning algorithm that shows high performance on visual imagery classification tasks are convolutional neural networks (CNN). A convolution is a mathematical operation of two functions that produces a third function to express

[5] A simple example of a set of features is using age, income and city to predict an individual's career.

how the shape of one function modifies the other (Goodfellow et al., 2016). Convolutional layers are stacked on top of each other, each passing their output to the next layer (Figure 24). In classification, this results in representations at deeper layers to contain information that is highly predictive of one condition or another (Goodfellow et al., 2016).

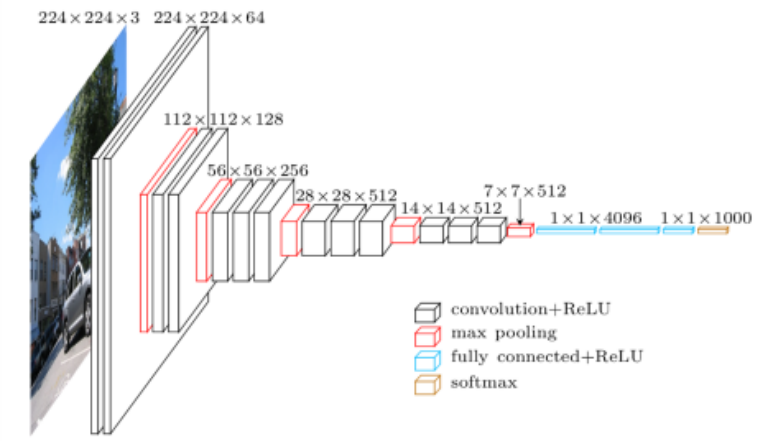

Figure 24. Schematic of a convolutional neural network applied to 2-dimensional images
This architecture, known as VGG, contains 19 convolutional layers (Simonyan & Zisserman, 2015).

Most recently, CNN architectures have been generalised to work on data represented in graphs instead of images. Deep graph convolutional neural networks (DGCNN) have been used to diagnose various neurological/psychiatric disorders using functional brain connectivity (Ktena et al., 2018; Parisot et al., 2018). Ktena et al. (2018) used a DGCNN to develop a distance metric between brain graphs of healthy controls and patients with Autism Spectrum Disorder (ASD) and found that the algorithm could accurately classify each condition with a high degree of accuracy. With the goal of applying these methods to DOC patients in Chapter 6, this chapter aims to develop and validate the use of DGCNN on healthy participants performing a volitional mental imagery task. As all of these participants are capable of performing the task, we will classify between participants either performing the task or resting, instead of using their resting state data as we will do in Chapter 6. We used the full brain connectivity graphs as well as individual canonical large-scale brain networks as inputs to the network. This work will establish a baseline for DGCNN performance, with the aim of applying this method to patient groups and comparing their performance with other machine learning methods. Futhermore, this study was

published as a conference paper at the 2nd Annual Cognitive Computational Neuroscience Conference.

5.2. Methods

5.2.1. Participants and Paradigm Specifications

Data was collected at the Wolfson Brain Imaging Centre at Addenbrooke's Hospital, Cambridge, UK. Ethics for this study were obtained from the Cambridgeshire 2 Regional Ethics committee. Twenty-five healthy participants were recruited for the study (19-57 years old, mean = 35.0, SD = 11.2, 13 males). Three participants were removed due to excessive head motion during scanning.

The behavioural paradigm consisted of five 30-second blocks where participants alternated between resting state and a motor imagery paradigm. During the motor imagery paradigm, participants were instructed to imagine being on a tennis court and swinging their arm to hit a tennis ball to an instructor on the opposite side of the net (Monti et al., 2010). The beginning of the motor imagery paradigm was cued with the word "Tennis" appearing on the screen and the beginning of the resting state paradigm began with the word "Rest" appearing on the screen. During the rest condition, participants were instructed to stay still, keep their eyes closed and think about nothing in particular.

5.2.2. Data Acquisition and Pre-processing

Data was acquired using a 3T Tim Trio Siemens system (Erlangen, Germany). We used a 12-channel head matrix transmit-receive coil with a fast-sparse 32 slice axial oblique sequence (TR = 2000ms, TE = 30ms, flip angle = 78°, voxel size = 3.0 x 3.0 x 3.0 mm, matrix size 64 x 64, field of view 192mm x 192mm, slice thickness = 3.0mm). We collected 150 EPI images in each subject's run. T-1 weighted MPRAGE high-resolution structural images were also acquired with 1.0 x 1.0 x 1.0 mm resolution (TR = 2250ms, TI = 900ms, TE = 2.99ms, flip angle = 9°).

Pre-processing was performed with Statistical Parametric Mapping 12 (SPM12; http://www.fil.ion.ucl.ac.uk/spm/) and MATLAB version 2017a

(http://www.mathworks.co.uk). The first five volumes for each were removed to eliminate saturation effects and achieve steady-state magnetisation. Following motion correction in the functional dataset, the participant's high-resolution structural images were coregistered to the mean EPI and segmented into grey matter, white matter, and cerebrospinal fluid masks. Next, the images were normalised to MNI space resampled to a resolution of 2 mm^3. Functional images were smoothed with a 6mm FWHM Gaussian kernel. To reduce residual movement-related and physiological artefacts, data underwent de-spiking with a hyperbolic tangent squashing function. Next, the aCompCor technique was used to remove the first 5 principal components of the signal from the white matter and cerebrospinal fluid masks, as well as 6 motion parameters and their first order temporal derivatives and a linear de-trending term (Behzadi et al., 2007). Functional images were then highpass filtered to remove low frequency fluctuations associated with scanner noise (0.009 Hz $< f$).

Because deep learning methods need a large amount of training data (Lecun et al., 2015), we augmented our data using dynamic functional connectivity. Each condition contained 75 volumes that were concatenated across the entire run. We used 50 volume-sliding windows across each condition to generate 25 correlation matrices per condition, per subject.

Functional connectivity was calculated by computing the correlation coefficient between time series from each of the 118 cortical regions from the Lausanne parcellation (Hagmann et al., 2008). As DGCNN cannot currently work with weighted edges, each correlation matrix was binarised keeping the top 20% of connections. We used the entire 118 region graph as well as 6 large-scale brain networks as input to the DGCNN. This allowed us to determine whether a specific large-scale network is most important for classifying cognitive state. These networks included the Aud, DMN, FPCN, SN, SM, and Vis. Node assignment to each network was calculated by overlapping each ROI with a network mask from Smith et al. (2009), as described in section 4.2.4.1 of Chapter 4.

5.2.3. Deep Graph Convolutional Neural Network

The DGCNN used here is adapted from Zhang et al. (2018) and consists of three sequential stages: 1) graph convolutional layers to extract node connectivity features; 2) a SortPooling layer to sort node features and equate input feature size, and 3) a series of classical convolutional and fully connected neural network layers to read the sorted graph representations and make predictions (Zhang, Cui, Neumann, & Chen, 2018). We use **A** to represent the adjacency matrix and *n* to denote the number of nodes. Each node has a *c*-dimensional vector where we use $\mathbf{X} \in \mathbb{R}^{nxc}$ to denote the feature matrix for the graph, with each row representing a node's feature vector. Each feature vector contains information about the node's neighbours and the node itself such as degree and connectivity strength with its neighbours (in the form of z scores). Each column, *c*, in **X** is known as a *feature channel*.

The graph convolutional layers aggregate node information from neighbouring nodes to extract multi-scale graph substructures important for classification. For each graph convolutional layer, given a graph **A** and its node information $\mathbf{X} \in \mathbb{R}^{nxc}$, information is propagated through the network with:

$$\mathbf{Z} = f(\widetilde{\mathbf{D}}^{-1}\,\widetilde{\mathbf{A}}\mathbf{X}\mathbf{W})$$

where $\widetilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ is the adjacency matrix plus the identity matrix. The inclusion of the identity matrix results in self-loops, which means that for the node undergoing convolution, it is multiplying its neighbours and its own feature vectors. $\widetilde{\mathbf{D}}^{-1}$ is the inverse of the diagonal degree matrix. We take the inverse of this matrix so as to normalise $\widetilde{\mathbf{A}}$, otherwise, $\widetilde{\mathbf{A}}$ will change the scale of the feature vectors. $\mathbf{W} \in \mathbb{R}^{cxc'}$ is a matrix of trainable weights and *f* is a nonlinear activation function.

The convolution itself contains four separate steps. The first step is a linear feature transformation applied to the node feature matrix by **XW**. This maps the *c* feature channels to the *c'* feature channels in the next layer and the trainable weights are shared amongst all nodes. The next step, $\widetilde{\mathbf{A}}\mathbf{Y}$, where $\mathbf{Y} \coloneqq \mathbf{XW}$, passes the neighbouring node's feature vectors (plus the self-loop) to the node that is undergoing convolution. To show this fully we see that:

$$(\widetilde{\mathbf{A}}\mathbf{Y})_i = \sum_j \widetilde{\mathbf{A}}_{ij}\mathbf{Y}_j = \mathbf{Y}_i \sum_{j\in\Gamma(i)} \mathbf{Y}_j$$

where the i^{th} row of the matrix is the summation of $\mathbf{Y}_i$ itself and $\mathbf{Y}_j$ from i's neighbouring nodes. Step three is the normalisation of row i by multiplication with $\widetilde{\mathbf{D}}_{ii}^{-1}$. The final step applies a nonlinear activation function f (in this case a rectified linear unit or ReLU), which outputs the final graph convolution result (Zhang et al., 2018). The preceding four steps aggregate *local* nodal information. To extract multi-scale substructures from the data, thus resulting in a richer featurisation and ultimately better classification, we stack multiple graph convolutional layers on top of one another (Figure 25). This is done using:

$$\mathbf{Z}^{t+1} = f(\widetilde{\mathbf{D}}^{-1}\widetilde{\mathbf{A}}\mathbf{Z}^t\mathbf{W}^t)$$

where $\mathbf{Z}^0 = \mathbf{X}$ and $\mathbf{Z}^t \in \mathbb{R}^{nxc_t}$ is the output of the t^{th} graph convolutional layer. Here, c_t is the number of output channels at layer t and the weight matrix $\mathbf{W}^t \in \mathbb{R}^{c_t \times c_{t+1}}$ maps c_t channels in the current layer to c_{t+1} channels in the following layer. The final graph convolutional layer concatenates the output by:

$$\mathbf{Z}^{1:h} := [\mathbf{Z}^1, \ldots, \mathbf{Z}^h]$$

where h is the number of graph convolutional layers. This final concatenated output represents a multi-scale feature descriptor for each node in the graph.

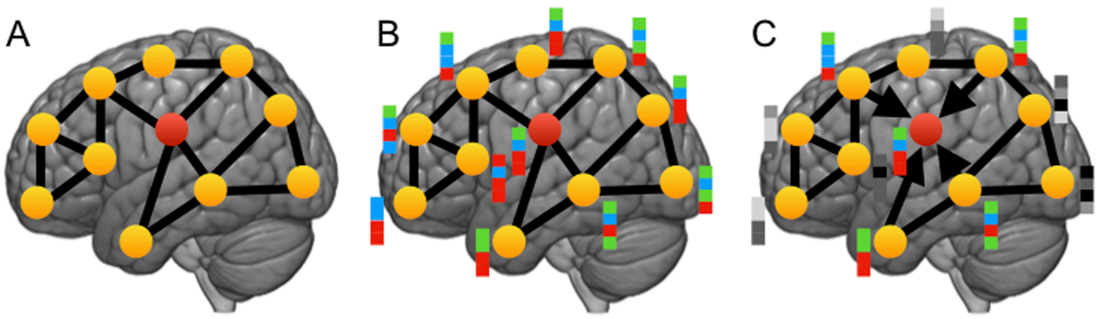

Figure 25. Schematic of the layer-wise propagation rule in a graph convolutional network
A shows a brain graph and highlights one node (in red) that will have the rule applied to it. **B** demonstrates that each node in the graph has its own feature vector associated with it, which contains basic information like connectivity strength with its neighbours and which large-scale brain network it belongs to. **C** shows that the red node's neighbouring nodes will combine their feature vectors, a process that when applied to every node in the graph through multiple layers results in multi-scale feature descriptors that can then be used to classify which condition this data is in.

The SortPooling layer sorts the output of the graph convolutional layer so that node feature vectors are pooled together and outputted in a consistent order. This is important because the final 1D convolutional step is most effective at classification when features are presented in a consistent order (Zhang et al., 2018). This final step consists of two layers of classical 1D convolutional layers, each with a convolutional layer, followed by a rectified linear unit activation function and maxpooling layer. This is followed by a fully connected layer and a softmax layer for classification. Each classification analysis was trained for 200 epochs with a learning rate of 0.0001 and three convolutional layers. The complete dynamic functional connectivity data from 4 subjects were randomly allocated to a test set so as to not over-fit the classifier during training.

To assess the results of each binary classification we used the following metrics: Precision (True positives/True Positives + False Positives), recall (True Positives/True Positives + False Negatives), F1 Score (Harmonic mean of Precision and Recall), and the area under the receiver operating characteristic curve (ROC-AUC; Bishop, 2006). The ROC curve is a plot of True Positive Rate against False Positive Rate for different cut-offs of a diagnostic test and is a measure of trade-off between sensitivity (true positive rate) and specificity (1 – false positive rate). As our analysis has balanced classes (i.e. an equal number of examples for each cognitive state), the ROC-AUC was considered the most important metric (Davis & Goadrich, 2006). With unbalanced data, the true positive rate and false positive rate can reflect high accuracy but is largely uninformative. For example, in cancer diagnosis, there are typically very few cancerous cases, meaning that a classifier could simply predict all data points as non-cancerous and achieve a relatively high accuracy (Davis & Goadrich, 2006). In these cases, other metrics like the precision-recall curve are more informative. However, as we have a balanced dataset, the ROC curve is the most appropriate metric to use. A comprehensive list of results can be found in Table 1.

5.3. Results

We used a DGCNN to determine whether a specific large-scale brain network was most predictive of cognitive state in healthy participants performing a mental imagery

task. We first examined whether we could accurately classify between tennis imagery and resting states using a full cortical graph. This analysis resulted in a ROC-AUC of 0.595 on held out test data. We then used each of the large-scale networks as input and found that FPCN (ROC-AUC = 0.885) and DMN (ROC-AUC = 0.793) had much higher ROC-AUC scores than the other networks, suggesting they alter their connectivity pattern the most between the mental imagery and resting state conditions (Figure 26). We also combined these two networks, including both within and between network connections as input. This, however, resulted in a slightly poorer classification with a ROC-AUC of 0.75. The Aud had the next highest ROC-AUC at 0.667, followed by the SN (0.562), SM (0.48) and Vis (0.475).

	Precision	Recall	F1	ROC AUC
Auditory	0.64	0.635	0.637	0.667
Somatomotor	0.54	0.525	0.532	0.487
Visual	0.51	0.525	0.517	0.475
Salience	0.54	0.545	0.542	0.562
DMN	0.7	0.7	0.700	0.793
FPCN	0.81	0.805	0.807	0.885
FPCN-DMN	0.7	0.7	0.7	0.75
Full Network	0.75	0.76	0.755	0.595

Table 1. Classification results for held out test data for each network input.

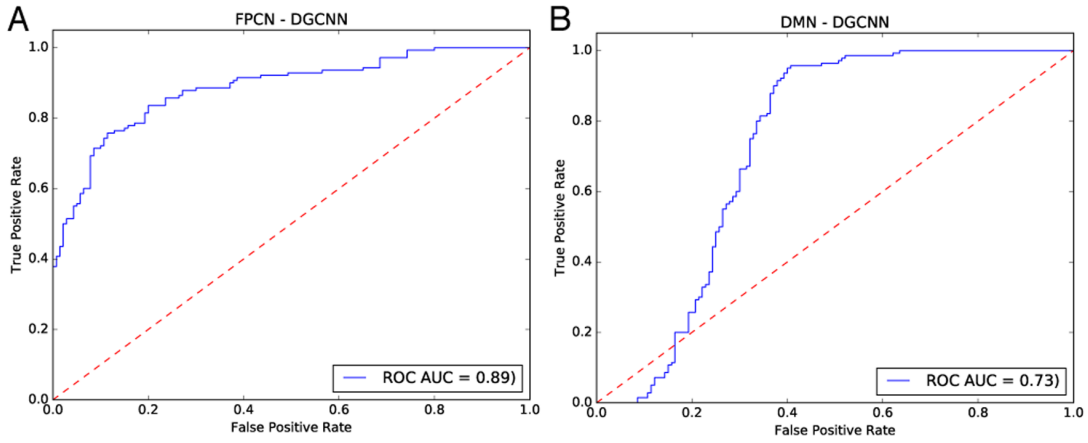

Figure 26. ROC Curve for FPCN and DMN
A shows ROC curve for FPCN and **B** shows the ROC curve for DMN. These show False Positive Rate on the X axis and True Positive Rate on the Y axis. Curves approaching the upper left-hand corner of the plot suggest better classification.

5.4. Discussion

The present study used functional connectivity fMRI, a mental imagery task, and a DGCNN to determine which large-scale brain networks were most predictive of mental imagery state.

We first found that the full cortical graph could predict mental imagery state with modest accuracy. Interestingly, the full graph had a relatively high precision and recall, but a relatively low ROC-AUC. This is likely due to the fact that ROC-AUC is calculated using prediction probabilities (values ranging between 0 and 1) and not binary classification measure (either 0 or 1). Essentially, this means that the full graph classification was often correct, however, the DGCNN was less confident in its decision suggesting that using specific brain networks as input may provide a better starting point for the DGCNN.

We also examined the predictive capacity of several large-scale cortical brain networks. We found that FPCN and DMN were most predictive of whether a participant was performing the tennis mental imagery task or in a resting block. We also found that using both within and between DMN and FPCN connections as input features resulted in a relatively high ROC-AUC, but did not improve the performance above using within network connectivity from each network. Previous work has shown that these networks dynamically interact during mental imagery tasks involving future planning (Beaty et al., 2018; Gerlach, Spreng, Madore, & Schacter, 2014). Gerlach et al. (2014) examined network connectivity in participants performing process simulation tasks. Specifically, participants were asked to think of goals, like eating healthier food or learning to paint and to plan the steps necessary in order to complete these goals. Functional connectivity analysis found that sub-regions of both DMN and FPCN, including the PCC, dorsolateral prefrontal cortex, and anterior inferior parietal lobule all correlated with task performance. This is somewhat in contrast to our results, in that between FPCN-DMN connectivity did not improve classification. This is likely due to the fact that our task involved switching between two different states, as opposed to continuously performing the mental imagery task.

Another earlier study showed that the typically anticorrelated DMN and SN are modulated by FPCN, which flexibly couples with either network depending on whether the task involves internally or externally directed attention (Spreng et al., 2013). As these are internally directed tasks, it is sensible that we found DMN, and not SN as being important for classifying between conditions. Our findings are therefore in line with these results in that they show FPCN and DMN are most predictive of a mental imagery task involving the switching between two states where attention is directed internally.

The present study has focused only on healthy participants. However, mental imagery tasks similar to the one used here have been used in the diagnosis and prognosis of patients with DOC (Monti et al., 2010; Owen et al., 2006). This work provides a baseline for Chapter 6, which aims to automatically detect awareness in these patients. Chapter 6 also compares the performance of the DGCNN with other machine learning classifiers to determine whether this more complex method provides greater discriminative ability between patient groups.

Chapter 6. Classifying volitional capacity in patients with disorders of consciousness

6.1. Introduction

DOC are a severe type of neuropathology characterised by perturbations to the neural systems underlying awareness and wakefulness (Giacino et al., 2014). As discussed in Chapter 1, these clinical states are often the result of relatively focal traumatic brain injury or more global injury induced by hypoxia. Diagnosing these patients is primarily done at the bedside by observing their motor response to a variety of commands (Giacino et al., 2014). However, accurately diagnosing patients is complicated by the fact that their behaviour can fluctuate throughout the day and the possibility that they suffer from additional cognitive or sensory impairments. These challenges have motivated the development of neuroimaging biomarkers that can aid in the diagnosis and prognosis of patients with DOC.

PET and fMRI have also been extensively examined as methods for assisting in the diagnosis of patients with DOC. Studies using these methods either use activation or connectivity measures during resting state (Boly et al., 2009; Demertzi et al., 2015; Di Perri et al., 2016), sensory stimulation (Coleman et al., 2007; Schiff et al., 2005), or active mental imagery paradigms (Bardin et al., 2012; Menon et al., 1998; Monti et al., 2010; Owen et al., 2006). Mental imagery paradigms are of particular interest, as this work has shown that a subset of patients can covertly respond to commands while in the scanner. Monti et al. (2010) showed that 5 of 54 patients reliably modulated their brain activity when instructed to imagine themselves playing tennis (motor response) or imagine walking around their house (navigation response). Bardin et al. (2012) expanded on this work by using a more advanced statistical analysis, multivariate pattern statistical analysis to examine distributed patterns of responses while patients were instructed to imagine performing a motor task. They found that 3 of the 5 patients who participated in their study had responses that could reliably be classified as task or rest. While these approaches demonstrate the utility of using mental imagery tasks to assess volitional ability in patients, they are limited by the fact that

many institutions lack the facilities and technical expertise to implement these paradigms and their associated analyses. Therefore, utilising resting state measures would provide a more generalizable clinical tool for the assessment of this patient group.

Demertzi et al. (2015) combined resting state functional connectivity with a machine learning algorithm to classify patients with different diagnoses. Specifically, they quantified connectivity measures in 6 large-scale brain networks and a support vector machine (SVM) classifier and found that each network had a high degree of discriminative capacity (greater than 80%) when classifying patients as MCS or UWS. This study demonstrated that functional connectivity measures have significant discriminatory power when assessing DOC patients. More recently, Di Perri et al. (2016) demonstrated that large-scale brain network interactions, specifically the interaction between DMN and the putative "task-positive" network with which it is anticorrelated (Fox et al., 2005), also predicts the level of consciousness in patients with DOC. This evidence suggests that the interaction between these specific networks may be a useful feature set for a machine learning algorithm.

Taken together, a useful clinical tool for deepening our understanding of DOC patients would provide information on the patient's internal state using a simple measurement like resting state fMRI. Ideally, this tool would not be a black box, but would have a degree of interpretability based in cognitive neuroscience. For this reason, the present study utilises resting state functional connectivity and deep learning, not to provide a specific diagnosis, but to classify patients with DOC as being able or unable to respond to volitional mental imagery tasks. This decision is motivated by the fact that a clinical diagnosis will always be accompanied with extensive behavioural evaluation, and therefore stratifying patients by volitional ability will provide clinicians with more actionable information to aid a diagnosis, and not simply reinforce their current assessment. Resting state fMRI data was collected from patients who have been shown to be able (VC group) or unable (non-VC group)

to respond to a volitional mental imagery task. We then applied functional connectivity analyses to determine if there were significant differences in network dynamics between VC and non-VC patients, providing us some grounding in basic cognitive network neuroscience. Finally, we implemented a cutting-edge deep graph convolutional neural network (DGCNN) to automatically differentiate between groups and to determine whether a specific brain network provided the most discriminative capacity between patients capable of performing the volitional mental imagery task.

6.2. Methods

6.2.1. Patients

A sample of 71 DOC patients were included in this study. Patients were recruited from specialised long-term care centres. To be invited to the study, patients must have had a DOC diagnosis, consent from their legal representative, and were capable of being transported to Addenbrooke's Hospital. The exclusion criteria included any medical condition that made it unsafe for the patient to participate (decision made by clinical personnel blinded to the specific aims of the study or any reason they are unsuitable to enter the MRI scanner environment (e.g. non-MRI-safe implants)). After admission, each patient underwent clinical and neuroimaging testing. Patients spent a total of 5 days (including arrival and departure days) at Addenbrooke's Hospital. Medication prescribed to each patient was maintained during scanning.

As we are interested in examining specific graph metrics, patients were systematically selected for inclusion in the final cohort based on the following: 1) large focal brain damage (i.e. more than 1/3 of one hemisphere) as stated by an expert in neuroanatomy blinded to the patients' diagnoses; 2) excessive head motion during resting state scanning (i.e. greater than 3mm in translation and/or 3 degrees in rotation); 3) suboptimal segmentation and normalization of images. This process left us with a total of 30 patients (VS = 12, MCS = 18). The positive responder group had 3

VS patients and 10 MCS patients, two of which were eMCS. The negative responder group had 9 VS patients and 8 MCS patients, none of which were eMCS. Patients had an age range of 19-70 (mean = 38.2, std. dev. = 15.8; mean positive responders = 37.5, std. dev. = 15.6; mean negative responders = 38.8, std. dev. = 16.5).

6.2.2. Mental Imagery Tasks

The mental imagery tasks used here have been previously used to assess awareness in DOC patients (Owen et al., 2006; Monti et al., 2010; Bardin et al., 2012) and have been validated in healthy participants (Fernández-Espejo, Norton, & Owen, 2014). Briefly, each patient was asked to perform two mental imagery tasks. In the motor imagery task, the patients were instructed to imagine being on a tennis court swinging their arm to hit the ball back and forth with an imagined opponent. The spatial imagery task involved the patient imagining walking to each room of their home or navigating the streets of a familiar city and to picture what they would observe if they were there. Each task consisted of five rest-imagery block cycles (30 seconds each). Each rest block was cued with the spoken word "relax" and each imagery block was cued with the spoken word "tennis" or "navigation". During the resting block, patients were instructed to stay still and keep their eyes closed.

6.2.2. Data Acquisition

Data was acquired at Addenbrooke's Hospital in Cambridge, UK, on a 3T Tim Trio Siemens system (Erlangen Germany). T1-weighted images were acquired with an MP-RAGE sequence (TR = 2300ms, TE = 2.47ms, 150 slices, 1 x 1 x 1.2mm resolution). Functional images were acquired using an echo planar sequence (TR = 2000ms, TE = 30ms, flip angle = 78°, 32 slices, 3 x 3 x 3.75mm resolution).

The diffusion-weighted dataset used an echo planar sequence (TR = 8300 ms, TE = 98 ms, matrix size = 96 x 96, 63 slices, slice thickness = 2 mm, no gap, flip angle = 90°). This included diffusion sensitising gradients applied along 12 non-collinear directions with 5 b values that ranged from 340 to 1590 s/mm^2 and 5 b = 0 images. Using multiple b-values increases the accuracy of DTI results (Correia, Carpenter, & Williams, 2009).

6.2.3. Data Pre-processing

6.2.3.1. Brain Network Parcellation

For consistency and comparability with previous chapters, the following analyses used the same brain parcellation. We used the 129 region Lausanne Parcellation (Hagmann et al., 2008) which consists of 64 regions per hemisphere and 1 bilateral brainstem region. For large-scale brain networks we used the 7 masks from Smith et al., 2009). We used an in house-script built in MATLAB to calculate for each ROI the overlap with each sub-network to assign each ROI to one of these networks. The maximum overlap was used to assign the ROI to the respective sub-network. As previous work has shown that interactions between high-level cognitive networks are important in stratifying DOC patients (Di Perri et al., 2016), we only included the regions from the DMN, FPCN, and SN. This resulted in a total of 71 ROIs in our final analysis.

6.2.3.2. High-Resolution T1 and fMRI Data Pre-processing

Pre-processing was performed with Statistical Parametric Mapping 12 (SPM12; http://www.fil.ion.ucl.ac.uk/spm/) and MATLAB version 2017a (http://www.mathworks.co.uk/products/matlab/). The first five volumes for each were removed to eliminate saturation effects and achieve steady-state magnetization. High-resolution structural images were coregistered to the mean EPI and segmented into grey matter, white matter and cerebrospinal fluid masks (Ashburner & Friston, 2005). Next, the images were normalised to MNI space with a resolution of $2mm^3$. Functional images were smoothed with a 6mm FWHM Gaussian kernel. To reduce movement-related and physiological artefacts, data underwent de-spiking with a hyperbolic tangent -squashing function. Next, the aCompCor technique was used to remove the first 5 principal components of the signal from the white matter and cerebrospinal fluid masks, as well as 6 motion parameters and their first order temporal derivatives and a linear de-trending term (Behzadi et al., 2007). In line with the results obtained in Chapter 2, functional images were then highpass filtered to remove low frequency fluctuations associated with scanner noise (0.009 Hz < f).

6.2.3.4. Diffusion Tensor Imaging

Diffusion data was reconstructed using generalised q-sampling imaging (GQI) from DSI Studio (http://dsi-studio.labsolver.org/). For details on GQI, see Chapter 4.2.4.3. Default parameters from DSI Studio were used in the analysis. The selected diffusion sampling threshold was 1.25. A deterministic fibre-tracking algorithm was used (Yeh et al., 2013). The anisotropy ratio was 0.0518329 and the angular threshold was set to 60 degrees with a step size of 0:625 mm. Tracks with a length of less than 10mm were discarded. A 71 x 71 connectivity matrix was created by using each subject's parcellation and quantifying GFA between each region (Cohen-Adad et al., 2008; Tuch, 2004).

6.2.4. Statistical Analysis and Machine Learning

6.2.4.1. Univariate fMRI Analysis

Univariate fMRI analysis was conducted on all 30 patients for both the motor and spatial mental imagery tasks. The analyses were performed using FSL (version 5.0.9) (Jenkinson, Beckmann, Behrens, Woolrich, & Smith, 2012). The results of these analyses determined which patients would be placed in each classification condition. Patients who had patterns of activity associated with volitional capacity were placed in the Positive Responders condition, while patients that did not show these patterns were placed in the Negative Responders condition. For each functional scan, a general linear model consisting of contrasting periods of rest and active imagery was computed. Results were considered significant at a cluster level of $z > 2.3$ (corrected $p < 0.05$) (Monti et al. 2010).

6.2.4.2. Input Data for Machine Learning

The input data for each machine learning pipeline consisted of the functional connectivity, function-structure overlap as defined in Chapter 4, and graph theoretical measures used in Chapter 4 resulting in a total of 18 features per patient. These included functional connectivity (FC) for within DMN, FPCN, SN, between-network connectivity for DMN-FPCN, DMN-SN, and FPCN-SN, structure-function (FC-SC) relationships with DMN, FPCN, SN, and between DMN-FPCN, DMN-SN, System Segregation, modularity, local efficiency, global efficiency, provincial hubs, and

connector hubs. For more details on how each of these metrics is calculated, see sections 4.2.5 for functional connectivity, 4.2.6 for function-structure overlap, and 4.2.7 for graph theoretical metrics in Chapter 4. We also included CRS-R for each patient, totalling to 19 features per patient. These features were fed into each feature-based machine learning algorithm as a vector for each patient. Differences between classes were also tested for using a paired t-test. All significant p-values were Bonferroni corrected. Classes were defined as patients who positively responded to a volitional fMRI task (n = 13) and patients who did not respond to a volitional fMRI task (n = 17).

For the classical machine learning models, data were separated into training, and test sets. Test datasets are completely held out from classifier training, ensuring the classifier can generalise and does not over-fit the training data. 12 subjects were randomly separated into the test set (training set = 18) for the classical machine learning algorithm analyses. For the deep learning models, we split the dataset into a training (n = 24), validation (n = 3), and test set (n = 3). The validation dataset was necessary to perform Bayesian hyperparameter optimisation, explained below.

6.2.4.3. Machine Learning Models

Each machine learning algorithm was implemented with scikit-learn (Pedregosa, Weiss, & Brucher, 2011), a Python-based machine learning library.

6.2.4.3.1. Logistic Regression

Logisitic regression (LR) classification is a regression model with a categorical dependent variable. LR aims to fit a separating line between two classes that is a linear function of feature inputs (Ryali, Supekar, Abrams, & Menon, 2010). In grid search optimisation L2 regularisation with the following values for C [0.001, 0.01, 0.1, 1.0, 10 100, 1000] were tested. L2 regularisation is known as the least squared error and minimizes the sum of squared differences between the target value and estimated value. C is the inverse of regularisation strength, where smaller values result in stronger regularisation (Pedregosa et al., 2011).

6.2.4.3.2. k-Nearest Neighbours

k-Nearest Neighbours (KNN) is a nonparametric classification method. The input is the k closest training examples in multidimensional feature space. The output of the algorithm is the predicted class, which is based on a majority vote of its k nearest neighbours (Altman, 1992). For grid search optimisation, we tested a range of k neighbours from 3 to 10.

6.2.4.3.3. Support Vector Classification

Support Vector Machines (SVM) construct a hyperplane in high dimensional space that separates classes of input data. In grid search optimisation, both linear and rbf kernels were tested. The linear kernel attempts to fit a linear hyperplane between the data classes while the rbf kernel attempts to fit a Gaussian hyperplane between the data classes (Chang & Lin, 2011). For the linear kernel, we tested the following values for C [1, 10, 100, 1000]. For the rbf kernel, we test the following parameters for Gamma [1e-5, 1e-4, 1e-3, 0.01, 0.1, 0.2, 0.5, 1.0, 2.0] and the following parameters for C [1, 10, 100, 1000]. C is the regularisation strength and Gamma is a coefficient for the rbf kernel.

6.2.4.4. Sequential Feature Selection

Sequential feature selection comprises greedy search algorithms that reduce the feature dimension space. Specifically, they reduce an initial d-dimensional feature space to a k-dimensional feature subspace where k < d. The aim is to produce a subset of features that are most relevant to the particular classification problem. Here we use sequential forward selection (SFS), an algorithm that iteratively adds features starting with the most predictive, to a feature subset and runs cross-validation with the chosen classification algorithm. The minimum number of features to achieve the highest cross-validated accuracy is chosen as the final feature subset. It is important to note that we only used training data for this process and that test data were completely held out.

6.2.4.5. Grid Search Optimisation

Grid search optimisation conducts an exhaustive search over a number of specified parameters for a particular classifier. The parameter combination with the highest cross-validated accuracy is then used for training and ultimately prediction of test data. The input to this step is the subset of features from the previous SFS algorithm.

6.2.4.6. Deep Graph Convolutional Neural Network

The DGCNN used in this analysis was adapted from Zhang et al., 2018 (https://github.com/muhanzhang/DGCNN) and consists of three stages: 1) graph convolutional layers to extract node connectivity features; 2) a SortPooling layer to sort the node features and equate input features size for; and 3) a series of classical convolutional and fully connected neural network layers to read the sorted graph representations and make final predictions (Zhang et al., 2018).

Details of the DGCNN are described in section 5.2.3. Briefly, graph convolutional layers aggregate node information from neighbouring nodes to extract multi-scale graph substructures important for classification. Each graph convolutional layer includes $\boldsymbol{A}$, an adjacency matrix, $\boldsymbol{D}$, a diagonal degree matrix, $\boldsymbol{X}$, a node information matrix with nodes in the rows and c node features in the columns, and $\boldsymbol{W}$, a matrix of trainable parameters. Each layer contains four steps. First, a linear transformation is applied to the node information matrix $\boldsymbol{XW}$. This maps the c feature channels to c' features in the next layer. The second step propagates node information to neighbouring nodes. Step three normalizes each node's feature vector. The final step applies a nonlinear activation function and outputs the graph convolution.

The SortPooling layer sorts the output of the graph convolutional layer so that node feature vectors are pooled together and outputted in a consistent order. This is critical because the final 1D convolutional step is most effective at classification when features are presented in a consistent order. This final step consists of two layers of classical 1-dimensional convolutional layers, each with a convolutional layer, followed by a rectified linear unit (ReLu) activation function and maxpooling layer. A fully connected layer and a softmax layer for classification follow.

The complete dynamic functional connectivity data from 3 subjects were randomly allocated to a test set so as to not over-fit the classifier during training. As deep learning models take much more time to train than classical algorithms (Snoek et al., 2015), we used Bayesian hyperparameter optimisation in python using pyGPGO (http://pygpgo.readthedocs.io/en/latest/) as opposed to grid search to optimise our models. Briefly, Bayesian optimisation uses previous results from the model to compute the posterior expectation of the loss function and then uses this expectation to find the next optimal set of hyperparameters to improve the model (Snoek et al., 2015). We performed Bayesian optimisation by training our model many times with increasingly optimised hyperparameters and evaluating its accuracy on the validation dataset. The best set of hyperparameters for each model was then used to make predictions on the final held-out test dataset.

6.3. Results

6.3.1. Positive and Negative Responders to Mental Imagery Tasks

We found a total of 13 patients who had a positive response to at least one mental imagery tasks (Figure 27). Patient 20 responded positively to both the motor imagery and spatial imagery tasks. Ten patients had positive responses to the spatial imagery task, while 4 patients had positive responses to the motor imagery task.

6.3.2. DMN, FPCN and SN Functional Connectivity in DOC

Previous work has shown that functional connectivity between DMN and putative TPNs is altered in DOC patients (Di Perri et al., 2016). We sought to expand on this work by examining within and between DMN connectivity, as well as within and between FPCN and SN connectivity in patients with DOC.

We did not observe any significant differences for within DMN connectivity ($t_{28)} = 0.203$, $p = 0.84$; Figure 28A), FPCN connectivity ($t_{(28)} = -1.681$, $p = 0.104$; Figure 28B), or SN connectivity ($t_{(28)} = 0.732$, $p = 0.103$; Figure 28C) for patients who responded to volitional tasks compared to those who did not. We did not observe any significant differences in DMN-FPCN connectivity ($t_{(28)} = 0.482$, $p = 0.633$; Figure 28D), or DMN-SN ($t_{28)} = 1.941$, $p = 0.0624$; Figure 28E) for DOC patients with

positive and negative responses to volitional tasks. However, we did see significantly greater connectivity between FPCN and SN ($t_{(28)} = 3.347$, $p = 0.0023$, Bonferroni corrected $p = 0.042$; Figure 28F) in patients who responded to a volitional task compared to those who did not.

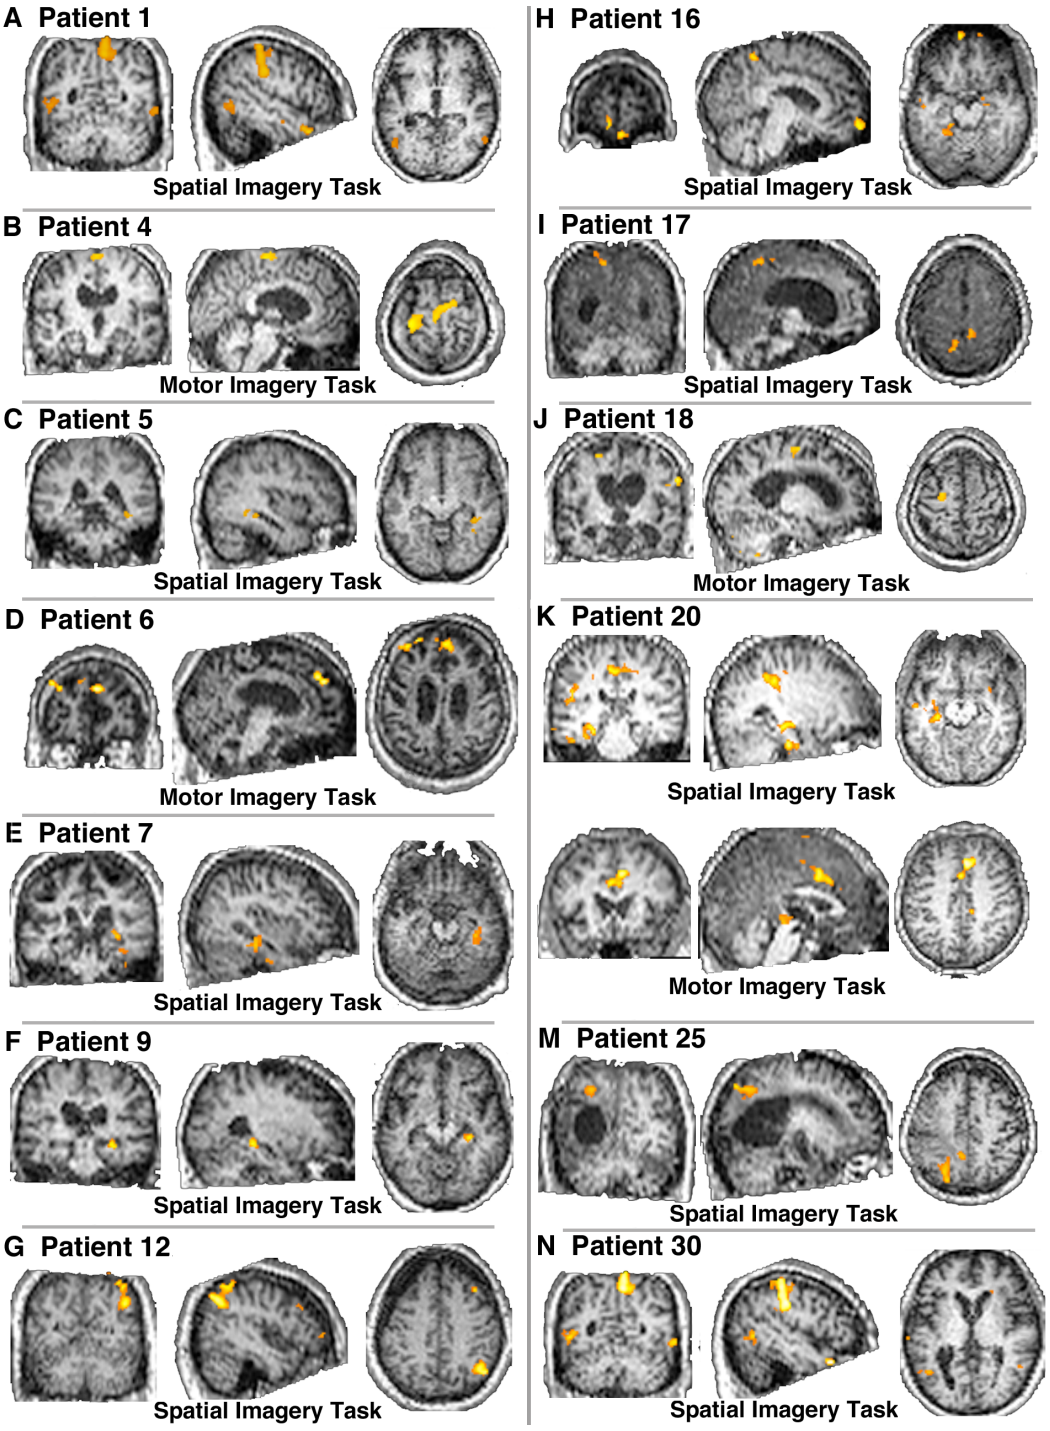

Figure 27. Positive responders to at least one mental imagery task
We found a total of 13 patients who responded to one or more mental imagery tasks. Patient 20 had positive responses to both the spatial and motor imagery task.

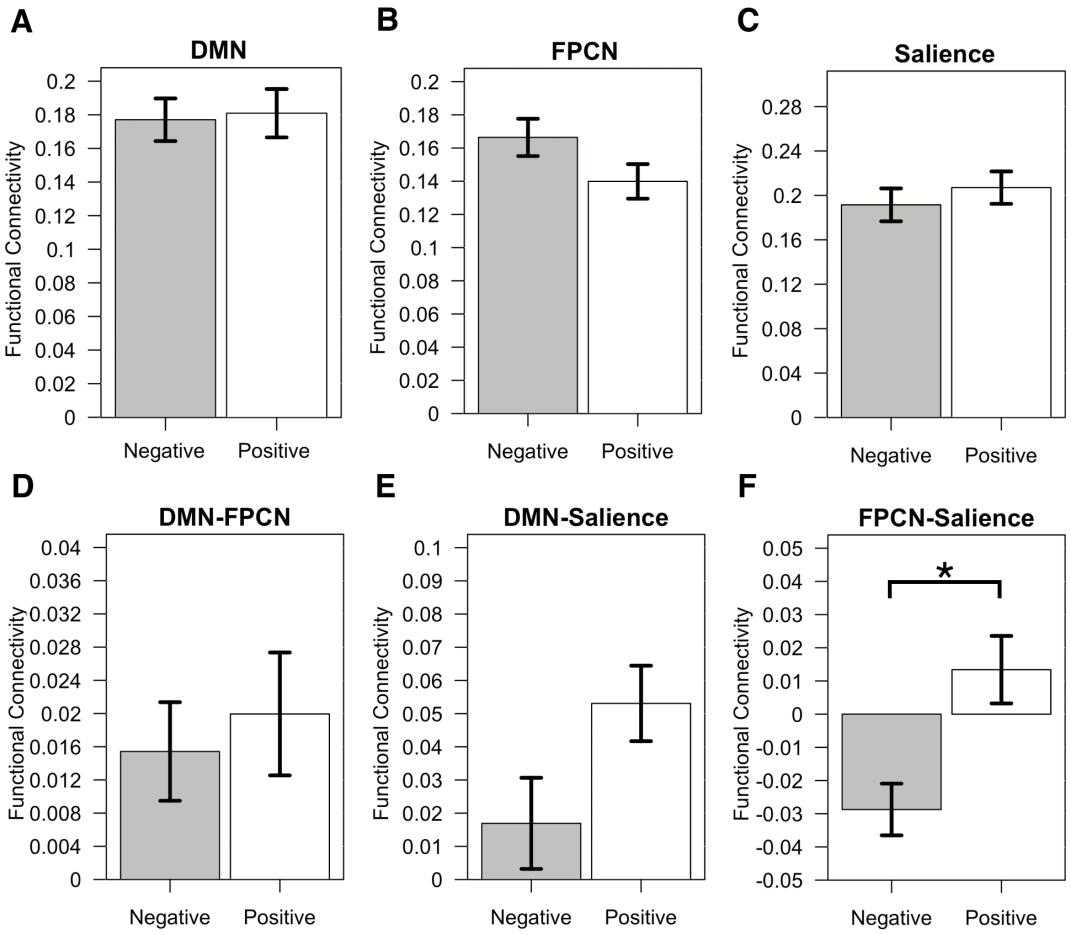

Figure 28. Functional connectivity within and between networks in DOC
A shows within DMN connectivity ($t_{(28)}$ = 0.203, p = 0.84), **B** shows within FPCN connectivity ($t_{(28)}$ = -1.681, p = 0.104), and **C** shows within SN connectivity ($t_{(28)}$ = 0.732, p = 0.103) between positive and negative responders to volitional tasks. **D** shows between DMN-FPCN connectivity ($t_{(28)}$ = 0.482, p = 0.633), **E** shows a between DMN-SN connectivity ($t_{28)}$ = 1.941, p = 0.0624) and **F** shows a significant difference between FPCN-SN connectivity ($t_{(28)}$ = 3.347, p = 0.0023, Bonferroni corrected p = 0.042) between positive and negative responders to volitional tasks.

6.3.3. DMN, FPCN and SN Function-Structure Overlap in DOC

We did not see any significant differences in function-structure overlap within DMN ($t_{(28)}$ = -0.615, p = 0.544; Figure 29A), within FPCN ($t_{(28)}$ = 0.793, p = 0.434; Figure 29B), or within SN ($t_{(28)}$ = 1.355, p = 0.186; Figure 29C). We also did not observe any differences in function-structure overlap between DMN-FPCN ($t_{(28)}$ = -0.761, p = 0.453; Figure 29D), DMN-SNs ($t_{(28)}$ = 1.889, p = 0.069; Figure 29E), or FPCN-SNs ($t_{(28)}$ = -0.408, p = 0.186; Figure 29F).

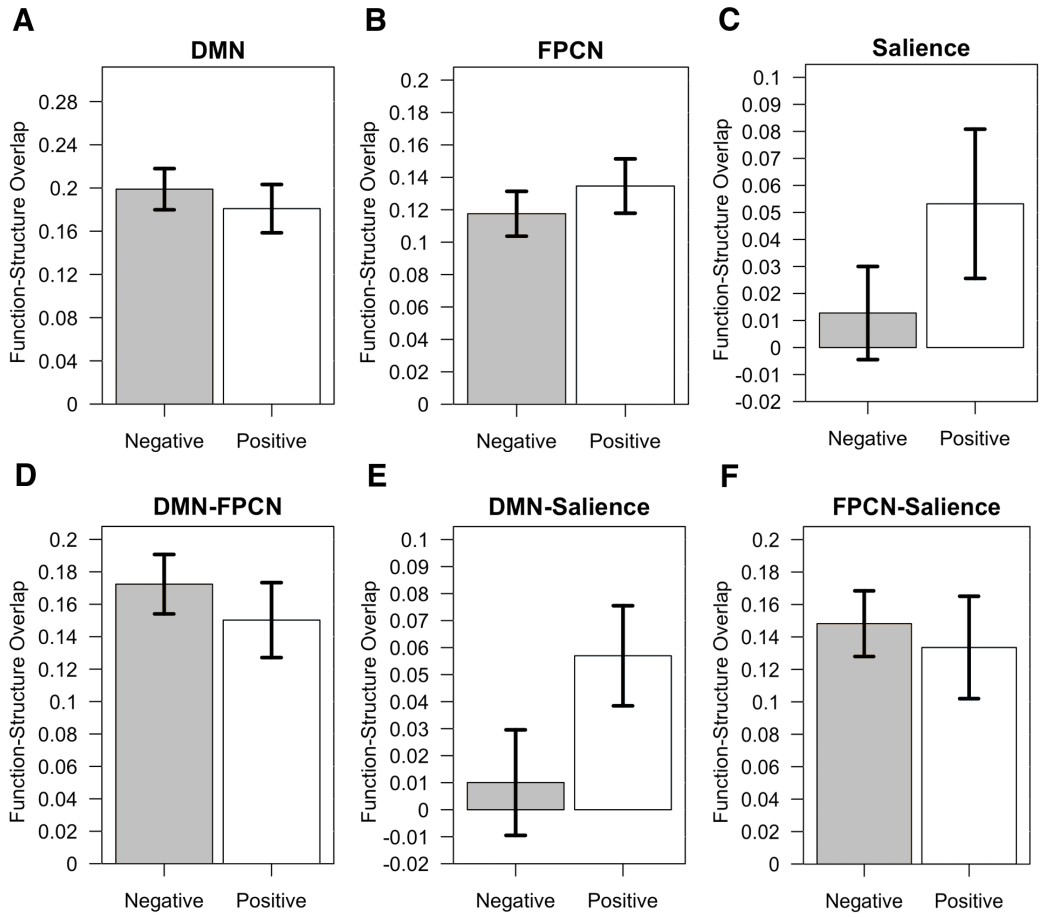

Figure 29. Function to structure overlap within and between networks in DOC
A shows within DMN function-structure overlap ($t_{(28)}$ = -0.615, p = 0.544), **B** shows within FPCN function-structure overlap ($t_{(28)}$ = 0.793, p = 0.434), and **C** shows within SN function-structure overlap ($t_{(28)}$ = 1.355, p = 0.186) between positive and negative responders to volitional tasks. **D** shows between DMN-FPCN function-structure overlap ($t_{(28)}$ = -0.761, p = 0.453), **E** shows between DMN-Salience function-structure overlap ($t_{(28)}$ = 1.889, p = 0.069) and **F** shows between FPCN-Salience function-structure overlap ($t_{(28)}$ = -0.408, p = 0.186) between positive and negative responders to volitional tasks.

6.3.4. Integration and segregation graph theoretical measures in DOC

We saw significantly greater system segregation in patients with a negative response compared to a positive response to the volitional tasks ($t_{(28)}$ = -3.524, p = 0.0018, Bonferroni corrected p = 0.0318; Figure 30A). We did not see any other significant differences in measures of network segregation in patients with DOC, including modularity ($t_{(28)}$ = -0.949, p = 0.35; Figure 30B), local efficiency ($t_{(28)}$ = 0.801, p = 0.43; Figure 30C), or the number of provincial hubs ($t_{(28)}$ = 0.55, p = 0.583; Figure 30D). We did not observe any significant differences in global efficiency ($t_{(28)}$ = -1.01, p = 0.321; Figure 30E) or in the number of connector hubs ($t_{(28)}$ = -1.12, p = 0.269; Figure 30F) between groups of patients that had either positive or negative responses to a volitional task.

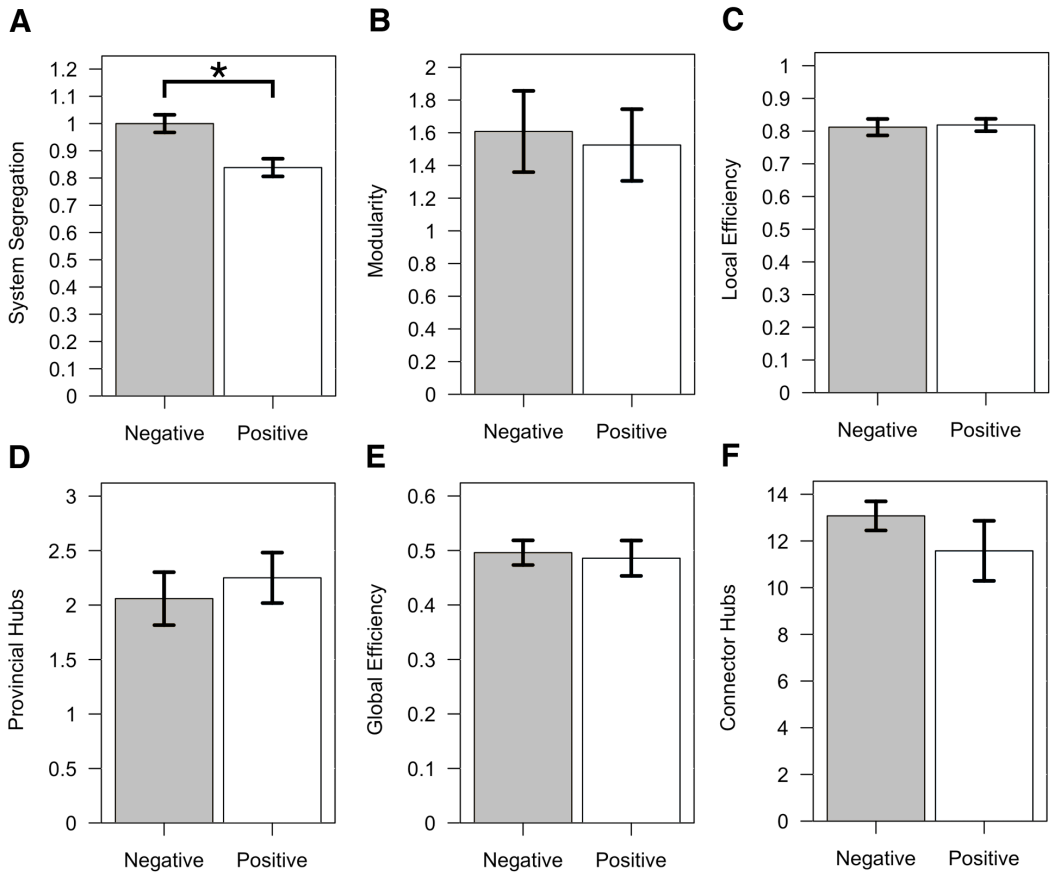

Figure 30. Graph theoretical measures in DOC patients
A shows a significant difference in system segregation (T(28) = -3.524, p = 0.0018, Bonferroni corrected p = 0.0318) between positive and negative responders to volitional tasks. **B** shows modularity (T(28) = -0.949, p = 0.35), **C** shows local efficiency (T(28) = 0.801, p = 0.43). **D** shows the number of provincial hubs (T(28) = 0.55, p = 0.583) between awake and moderate sedation conditions. **E** shows global efficiency (T(28) = -1.01, p = 0.321) and **F** shows the number of connector hubs (T(28) = -1.12, p = 0.269) between positive and negative responders to volitional tasks.

6.3.5. Feature Selection

For SFS in the DOC data, LR had a maximum feature accuracy of 0.96, meaning that as features were sequentially added to the classification, this was the highest value achieved. The subset features were CRS-R, DMN-SN FC, FPCN-FPCN FC, FPN-SN FC, System Segregation, Local Efficiency, Number of Provincial Hubs, Number of Connector Hubs, DMN-FPCN FC-SC, FPCN-FPCN FC-SC, FPCN-SN FC-SC, and SN-SN FC-SC. For KNN, SFS resulted in a maximum feature accuracy of 0.86 and included Local Efficiency, Number of Provincial Hubs, Number of Connector Hubs, and DMN-SN FC-SC in the feature subset. For SVM, SFS showed a maximum feature accuracy of 0.9 and selected CRS-R, DMN-FPCN FC, DMN-SN FC, FPCN-SN FC, System Segregation, Local Efficiency, Number of Provincial Hubs, Number of

Connector Hubs, DMN-FPCN FC-SC, FPCN-SN FC-SC and SN-SN FC-SC as a final feature subset.

6.3.6. Classification Results

For classifying DOC patients as positive or negative responders to a volitional task, we found that LR had a precision of 0.7, recall of 0.67, f1-score of 0.68, maximum training accuracy of 1.0, test accuracy of 0.66 and AUC of 0.67. KNN had a precision of 0.67, recall of 0.67, f1-score of 0.67, maximum training accuracy of 0.94, test accuracy of 0.66, and AUC of 0.62. SVM showed the best performance, with a precision of 0.74, recall of 0.74, f1-score of 0.74, maximum training accuracy of 0.88, test accuracy of 0.75, and AUC of 0.75. Comprehensive statistics can be found in Table 2.

Algorithm	Condition	Precision	Recall	f1-score	Test ACC	ROC-AUC
Logistic Regression	Negative	0.8	0.67	0.73		
	Positive	0.5	0.67	0.57		
	Avg/Total	0.7	0.67	0.68	0.66	0.67
KNN	Negative	0.75	0.75	0.75		
	Positive	0.5	0.5	0.5		
	Avg/Total	0.67	0.67	0.67	0.66	0.62
SVM	Negative	0.78	0.88	0.82		
	Positive	0.67	0.5	0.57		
	Avg/Total	0.74	0.74	0.74	0.75	0.72

Table 2. Classification Results for LR, KNN, and SVM
Classifying positive versus negative DOC responders to volitional tasks using classical machine learning algorithms on brain network data.

In addition to using these three classical machine learning algorithms, we also used a more advanced DGCNN architecture to classify patients. An advantage that convolutional neural networks have over other types of machine learning algorithms is that they require less input feature engineering since the network is typically capable of learning the appropriate features associated with each classification condition (Lecun et al., 2015). With this in mind, we used each patient's raw network values, as opposed to calculated functional connectivity measures, as input to our deep learning model. We found using the combined connectivity graph (meaning connections within and between DMN, FPCN and SN were included) for each network in our analysis as input resulted in a relatively high classification accuracy (test accuracy = 0.71; ROC-AUC = 0.79; Figure 31B). However, when focusing specifically on within-FPCN connectivity as input, the accuracy of classifying volitional ability in DOC patients substantially increased (test accuracy = 0.85; ROC-AUC = 0.92; Figure 31A). Using within-network connectivity from DMN reduced accuracy, however, it still performed relatively well (test accuracy = 0.62; ROC-AUC = 0.67; Figure 31C). Within-SN connectivity showed very poor classification accuracy (test accuracy = 0.32; ROC-AUC = 0.29; Figure 31D). For between-network connectivity, we found that using FPCN-SN internetwork connections somewhat reliably classified patients (test accuracy = 0.71; ROC-AUC = 0.82; Figure 31E). Both DMN-FPCN (test accuracy = 0.6; ROC-AUC = 0.66; Figure 31F) and DMN-Salience (test accuracy = 0.55; ROC-AUC = 0.54; Figure 31G) performed worse than FPCN-SN connectivity. Comprehensive statistics can be found in Table 3.

Network	Condition	Precision	Recall	f1-score	Test Acc	ROC-AUC
All	Negative	0.97	0.59	0.74		
	Positive	0.54	0.96	0.69		
	Avg/Total	0.83	0.72	0.72	0.71	0.79
DMN	Negative	0.76	0.63	0.69		
	Positive	0.45	0.6	0.51		
	Avg/Total	0.65	0.62	0.63	0.62	0.67
FPCN	Negative	0.9	0.86	0.88		
	Positive	0.75	0.82	0.78		
	Avg/Total	0.85	0.85	0.85	0.85	0.92
Salience	Negative	0.49	0.32	0.39		
	Positive	0.19	0.32	0.24		
	Avg/Total	0.39	0.32	0.34	0.32	0.29
DMN-FPCN	Negative	0.75	0.61	0.67		
	Positive	0.43	0.6	0.5		
	Avg/Total	0.64	0.6	0.62	0.6	0.66
DMN-Sal	Negative	0.68	0.62	0.65		
	Positive	0.35	0.41	0.38		
	Avg/Total	0.57	0.55	0.56	0.55	0.54
FPCN-Sal	Negative	0.89	0.64	0.75		
	Positive	0.54	0.85	0.66		
	Avg/Total	0.78	0.71	0.72	0.71	0.82

Table 3. Classification Results for DGCNN Across Brain Networks
Classifying positive versus negative DOC responders to volitional tasks using DGCNN with whole-brain and network-specific input features.

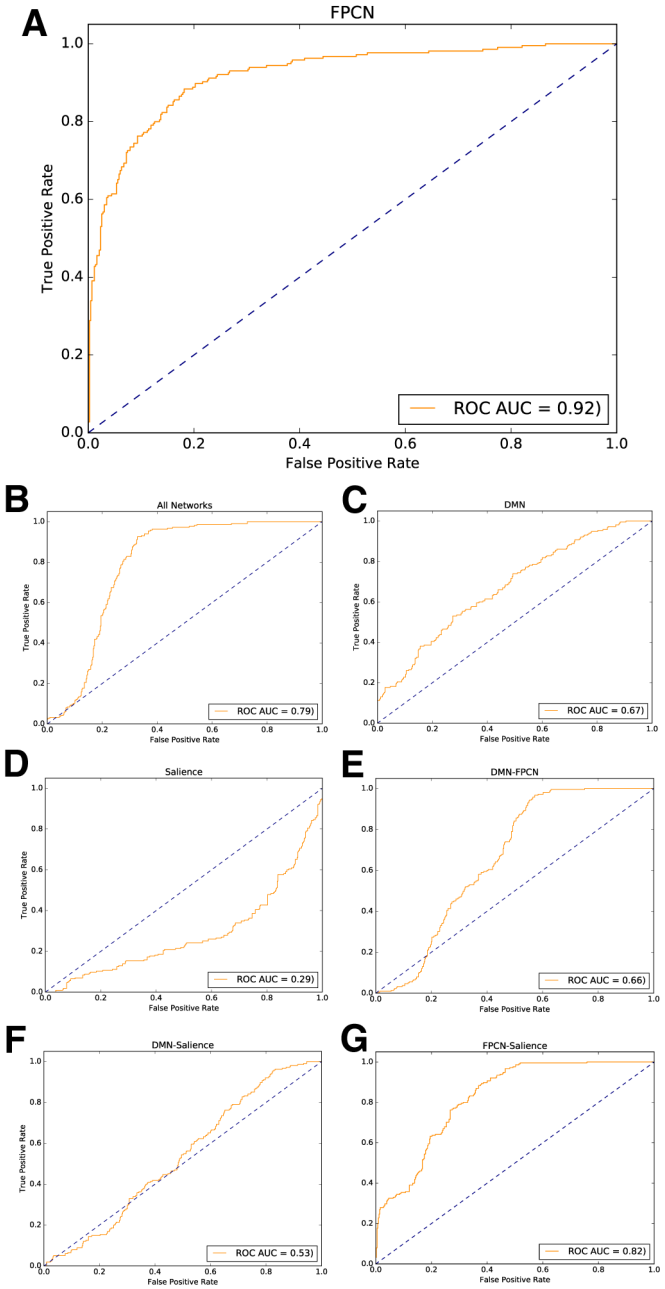

Figure 31. DGCNN ROC curves for classifying volitional ability in DOC patients
These graphs represent the receiver operator characteristic curves for each condition and each classification algorithm. The X-axis shows the False Positive Rate (or 1 – specificity), and the Y-axis shows the True Positive Rate (or sensitivity) of each classifier. These graphs represent the relative trade-off between true positive (benefits) and false positives (costs) for each classifier. The area under the curve (AUC) metric characterises the probability of the classifier ranking a randomly chosen positive example higher than a randomly chosen negative example. **A** shows the ROC curve for FPCN (AUC = 0.92). **B** shows the ROC curve for using all networks as input (AUC = 0.79). **C** shows the ROC curve for DMN (0.67). **D** shows the ROC curve for SN (AUC = 0.29). **E** shows the ROC curve for the interaction between DMN and FPCN (AUC = 0.68). **F** shows the ROC curve for DMN-SN interactions (AUC = 0.53). G shows the ROC curve for FPCN-SN interactions (AUC = 0.82).

6.4. Discussion

The present study was aimed at developing a resting state functional connectivity-based machine learning classifier to distinguish between patients with DOC who are either able or unable to respond to a volitional task. This builds upon previous work in the field by 1) using volitional ability, instead of clinical diagnosis, as classification conditions, 2) using a novel data augmentation approach, extensive hyperparameter optimisation, and several machine learning algorithms to determine which method best fits the data, and 3) identifying the importance of the FPCN in differentiating between patients with preserved volitional capacity. Additionally, we examined whether interactions within and between specific large-scale brain networks previously shown to be altered in DOC provide greater discriminatory power. We found that classical machine learning algorithms, including logistic regression, KNN, and SVM, performed relatively well in distinguishing between patients able or unable to respond to volitional tasks, with SVM achieving the highest score of the three (75% accuracy). We then used a more advanced algorithm, a DGCNN, and found that using only the FPCN as input resulted in an even higher accuracy (85% accuracy). These results show that it is possible to distinguish between DOC patients who are able or unable to respond to a volitional task using resting state fMRI and demonstrate its potential for clinical use.

By using a state-of-the-art DGCNN and comprehensive processing of our input features, we were able to achieve high classification scores for differentiating patients with and without volitional ability. Behavioural tasks have previously been used as a measure of covert awareness in DOC patients (Owen et al., 2006; Monti et al., 2010; Bardin et al., 2012; Naci, Graham, Owen, & Weijer, 2016). A seminal study by Owen et al. (2006) identified a patient who fulfilled all the clinical criteria for UWS who was able to respond to commands by modulating their fMRI activity using the same volitional tasks used in the current study. This was followed by Monti et al., (2010) who found that of 23 patients diagnosed with UWS, 4 showed signs of covert awareness based on their responses to these tasks. Subsequently, several studies used similar approaches with various motor tasks to identify covert awareness in patients (Cruse et al., 2011; Bardin et al., 2012). We extend these findings by using several

features derived from resting state functional connectivity, as opposed to univariate or multivariate analyses of task-based fMRI. Previous work in applying machine learning to DOC patient data has primarily focused on stratifying patients based on clinical diagnosis. Demertzi et al. (2015) used a support vector machine and resting state functional connectivity to classify DOC patients and found that the auditory network and DMN provided the best discriminative capacity between patients and healthy controls. An additional study by the same group found that the auditory network specifically was best when differentiating between patients with UWS and MCS (Demertzi et al., 2014). The present study had very similar accuracy results when discriminating between patients capable or not of responding to mental imagery tasks. Despite these interesting results, we chose not to include auditory network in our analysis because it is not typically implicated in mental imagery tasks and because later work from the same group suggested that the interaction between high-level cognitive brain networks has a high discriminative capacity when stratifying DOC patients (Di Perri et al., 2016; discussed below).

Using several complementary approaches, we found that connectivity both within and between the FPCN was best at differentiating DOC patients with and without volitional ability. First, we found significantly greater connectivity between FPCN and SN in patients capable of performing at least one mental imagery task compared to patients that were not. We also found that within FPCN connectivity and FPCN-SN connectivity served as the two best input features for the DGCNN. The use of functional connectivity to characterise group-level changes in brain network dynamics in DOC patients has been explored extensively. Previous work suggests that the DMN has significantly reduced connectivity at deeper levels of unconsciousness (Boly et al., 2009; Demertzi et al., 2014; Di Perri et al., 2016; Vanhaudenhuyse et al., 2010). Most recently, Di Perri et al. (2016) demonstrated that both within and between-DMN connectivity significantly correlated with the level of consciousness of DOC patients. Our analyses did not reveal any significant differences between positive and negative responders to tasks in either within or between-DMN connectivity, suggesting that this connectivity measure may reflect clinical diagnosis, but not volitional capacity. We did, however, observe significantly greater connectivity between FPCN and SNs in patients who positively responded to a volitional task

compared to patients who did not respond. Furthermore, we found that using FPCN as input to the DGCNN resulted in a high classification accuracy between patient groups. This result reflects many previous studies suggesting that FPCN is vital for executive function and cognitive control (Demertzi et al., 2014; Dosenbach et al., 2007; Naci et al., 2014). Naci et al. (2014) examined whether individuals share common neural processes using a plot-driven movie- watching task. They showed that one behaviourally non-responsive patient displayed significantly similar patterns to control patients in response to the task. Notably, the pattern of activity they observed was largely in frontoparietal regions.

It is well known that acquisition of data for patients with DOC is limited, and this is compounded by our selection criteria including the need for both task and resting state fMRI to be collected for each specific patient. Deep learning models are known to require large datasets for training and often require tuning to find the best set of hyperparameters to fit the data (Krizhevsky, Sutskever, & Hinton, 2012; Lecun et al., 2015). To overcome this limitation, we used a novel application of dynamic functional connectivity to augment the amount of data our deep learning model would be trained on. Instead of creating brain network graphs for a subject's entire resting state data, we used a sliding window approach that generated a graph for 80 consecutive functional volumes. This method generated 215 samples per patient, allowing us to effectively train our deep learning model. Importantly, in order to avoid over-fitting our model, we validated and tested it on patients who were completely held out from training. Additionally, we used Bayesian Optimisation to find hyperparameters that best fit a validation dataset for each network feature input (Lecun et al., 2015). Together, these methods allowed us to properly train and tune our deep learning model in an objective way, resulting in a high test accuracy for FPCN as an input feature.

In conclusion, the present study used resting-state functional connectivity and several machine learning algorithms to differentiate DOC patients who are capable from those who are incapable of responding to a volitional task during fMRI scanning. We found that a state-of-the-art DGCNN using FPCN as an input feature provided the best accuracy when classifying patients into these conditions. We also

found significantly greater connectivity between two task-positive brain networks, namely the FPCN and SN, in patients who could perform a volitional task. These results show that FPCN connectivity is particularly important in preserving volitional ability when performing mental imagery tasks and highlights the potential utility of deep learning in providing a clinical tool for neurologists in assessing patients with DOC.

Chapter 7. Discussion

7.1. Thesis Summary

The nature of consciousness has traditionally been considered a question for philosophers. However, with the tools of modern neuroscience, understanding consciousness as a neurobiological process has become increasingly possible (Crick & Koch, 1990; Koch et al., 2016). There are many approaches to understanding the neurobiology of consciousness, from cognitive and behavioural paradigms to examining how consciousness is altered in various clinical settings. The focus of this thesis has been to utilise two of these clinical settings, that is, changes of consciousness levels during propofol-induced sedation and as a result of DOC, as experimental models for understanding the changes in system-wide connectivity patterns associated with alterations in consciousness.

As brain network interactions were the primary measure of interest in this thesis, Chapter 2 focused on a pre-processing step important for accurately assessing these patterns. Specifically, we were interested in examining whether the relationship between large-scale brain networks at rest and during behavioural tasks changed as a function of how the data was pre-processed. Early work suggesting that the DMN was inactive during tasks initially motivated this study (Fox et al., 2005; Raichle et al., 2001). More recent studies have suggested DMN involvement in cognitive processing (Vatansever et al., 2015a, 2017). However, without a systematic investigation of DMN connectivity during tasks, its functional contribution to cognition cannot be fully elucidated. Resting state connectivity is thought to occur maximally at low frequencies (< 0.1 Hz), resulting in most studies removing high frequencies during data pre-processing. In contrast, subtractive task analyses include high frequencies, as these are thought to be relevant in task performance. An emerging line of research explores resting fMRI data at higher frequency bands, examining the possibility that functional connectivity is a multi-band phenomenon. We bridged these parallel lines of research by examining the functional relevance of high frequencies in the relationship between DMN and TPN at rest and, importantly, also during execution of behavioural tasks. Our findings revealed that the inclusion of high frequencies alters between network

connectivity resulting in reduced anticorrelation between DMN and task-positive networks. Critically, increased positive connectivity was observed only during tasks, suggesting a critical role for high frequency fluctuations in functional integration. Moreover, within-DMN connectivity correlated with task reaction time only when high frequencies were included. The presence of high frequencies suggests a more complex connectivity pattern between DMN and task-positive networks, warranting their inclusion in task-related functional connectivity studies.

With these results in mind, Chapters 3 and 4 used highpass filtering of fMRI data to examine functional connectivity during propofol-induced sedation. To provide a brief review, propofol is a commonly used anaesthetic drug that binds with high affinity to GABA receptors in the mammalian brain. Previous work on its anaesthetic action has characterised either the biochemistry underlying propofol binding or the changes in brain network dynamics during sedation. Despite these advances, no published study so far has focused on understanding how propofol action at the molecular level results in changes in brain network connectivity. To bridge this gap, we combined whole-brain microarray data of GABAergic markers with multivariate functional connectivity analysis to determine whether specific subtypes of GABAergic interneurons drive the alteration in brain network dynamics observed during propofol-induced sedation. To do this, we first used whole-brain human microarray data from the Allen Institute for Brain Science to generate distribution maps for genes that mark the primary GABAergic cortical interneurons (Parvalbumin, Somatostatin, 5-HT3a). Next, we identified which brain region's connectivity patterns were most impacted by propofol sedation. We then correlated these multimodal cortical patterns and found that regions with a high concentration of parvalbumin-expressing interneurons significantly overlapped with regions that disconnected during sedation. Together, these results suggest that propofol binding to parvalbumin interneurons is associated with altered brain connectivity and demonstrates the utility of leveraging multimodal data sets to address multi-scale questions in neuroscience.

Chapter 4 sought to address another unanswered question of how propofol administration alters brain connectivity, specifically in how functional connectivity

between networks relates to their underlying structural substrate. We found that connectivity between the DMN and TPNs like FPCN and SN increases during propofol administration. This is significant because, during normal wakefulness, these networks are typically weakly correlated or even anticorrelated. Previous research using non-human primate models has also shown that functional connectivity has greater overlap with structural connectivity during sedation compared to normal wakefulness. However, relationships between specific large-scale brain networks have not yet been established. Therefore, the study presented in Chapter 4 aimed to expand these findings by investigating the relationships between the DMN, FPCN, and SN during propofol-induced sedation as well as how their functional connectivity corresponds to underlying structural connectivity in healthy humans. We found that propofol sedation resulted in significant increases in functional connectivity and a decrease in the number of anticorrelations between DMN and SN. Furthermore, we found that the overlap between functional and structural connectivity also increased significantly during sedation between these networks. No differences were observed in functional connectivity or structure-function overlap for any other network relationships. Finally, we found that system segregation, a graph theoretical metric used to assess network partitioning, was significantly decreased during sedation, suggesting networks were more integrated. These findings propose that connectivity between DMN and SNs drives the effects of propofol-induced sedation and that this relationship relies on increased traffic on structural connections between these networks.

Though there are undoubtedly more questions to address regarding how anaesthetic drugs affect brain connectivity, the remainder of the thesis focused on another clinical model of unconsciousness, namely patients with DOC. Chapter 6 comprehensively examined various measures of functional and structural connectivity in DOC patients. Additionally, it used several machine learning algorithms aimed at classifying these patients into different categories. As one of these machine learning measures has not previously been applied to functional connectivity, Chapter 5 provided methodological details and a validation experiment in healthy participants before application to DOC patients. Specifically, we developed a graph convolutional network classifier that can distinguish between mental imagery states in healthy

subjects using only functional connectivity data. Furthermore, we examined whether specific large-scale brain networks were more predictive than others and found that FPCN and DMN were most predictive of whether a participant was performing a mental imagery task or resting. These results demonstrate that graph convolutional networks could be developed to aid in the detection of awareness in DOC patients and show that changes in connectivity patterns in FPCN and DMN underlie alterations in mental imagery.

Chapter 6 examined various functional connectivity measures and several machine learning algorithms to characterise brain network connectivity in patients who were capable of responding to a volitional task as defined by task-based fMRI. Due to the severity of the injury and the dramatically altered state of consciousness, accurately diagnosing patients is a particularly challenging clinical decision. This has resulted in many previous studies using brain imaging to aid in the diagnosis and to assess the volitional capacity of patients with DOC. Recent work has also focused on combining brain imaging with machine learning to automatically diagnose patients as being in a MCS or UWS. While this work shows promise, it may have limited clinical usefulness because patient data labelling is dependent upon previous diagnosis. The goal of this study was to attempt to overcome this problem by stratifying patients by their volitional capacity rather than by their clinical diagnosis. We first determined whether each patient could perform the volitional task and then applied several machine learning algorithms to resting state functional MRI data to classify whether they were able or unable to perform these tasks.

We performed task-based fMRI analysis while patients performed a motor imagery and spatial imagery volitional task. Next, we calculated several resting state functional connectivity metrics for input to three classical machine learning algorithms, namely logistic regression, k-nearest neighbours, and support vector classification. Finally, we used a DGCNN on brain network data to 1) determine whether the use of a deep learning model performed better than standard algorithms and 2) identify a specific brain network that most significantly differentiates patients with or without volitional capacity. We found that the DGCNN was best at discriminating between patients with volitional capacity. This result was established

when using only the FPCN as input to the model. In addition to the machine learning classifier results, functional connectivity was significantly greater between FPCN and SN in patients with volitional capacity. Brain network system segregation, a graph theoretical metric, was significantly greater in patients with volitional capacity. We did not observe any other significant differences in functional connectivity, structure-to-function coherence, or graph theoretical metrics between patient groups. The integrity of FPCN connectivity, a brain network well known to play a key role in executive functions and cognitive control, is essential for volitional capacity preservation in patients with DOC. Deep learning methods can be used to distinguish between patients with and without volitional capacity with a high degree of accuracy, potentially providing an additional tool to aid in diagnosis, prognosis, and potential communication with DOC patients.

7.2. Main Contributions

As discussed in Chapter 1, the NCC can be described as the minimal neural mechanisms jointly sufficient for any one conscious experience (Koch et al., 2016). The NCC is typically studied using either cognitive and behavioural paradigms in healthy participants or by studying clinical populations, such as patients with DOC, or healthy participants undergoing consciousness-altering drug administration. The use of cognitive and behavioural paradigms can be understood as examining changes to local states of consciousness, such that they pinpoint normal processing of perceptual experiences, including mental imagery, bodily sensations, occurring thought, and affective experiences (Bayne et al., 2016). The clinical models used to study consciousness are more aligned with examining global states, in that they do not differentiate between specific contents of subjective conscious experiences, but by an individual's overall conscious condition.

With these points in mind, the primary contributions of this thesis include deepening our understanding of the neurobiology of global states of consciousness using propofol sedation (Chapters 3 and 4) and expanding on the network neuroscience underlying local states of consciousness using healthy participants and a mental imagery paradigm (Chapter 5). Furthermore, Chapter 6 combined concepts

from both local and global states by using the mental imagery paradigm to differentiate between DOC patients with and without volitional capacity. Additionally, several methodological advances were developed to facilitate experimentation, including improvements of data pre-processing (Chapter 2), application of a new machine learning method to a clinical population (Chapter 6), and use of novel data featurisation and augmentation approach for improving machine learning results (Chapters 5 and 6).

One important goal of this thesis was to examine large-scale brain network connectivity in patients with DOC. The motivation for this is twofold. First, due to stratified degrees of perturbation to consciousness, DOC patients are commonly studied to deepen our understanding of global states of consciousness. Second, and more importantly, further characterisation of the neuropathology associated with DOC will aid in improving patient care. Chapter 6 aimed to address both of these points by combining functional connectivity measures with machine learning to differentiate between patients with and without volitional capacity. This is a significant contribution to the literature because the starting point is volitional capacity (as opposed to diagnosis; Demertzi et al., 2015) and by the use of a novel machine learning method specifically designed for graphs, which was developed in healthy participants in Chapter 5. A significant finding from Chapters 5 and 6 is the importance of the FPCN in preserving volitional capacity. This is in line with previous work demonstrating the importance of the FPCN in mediating shifts between internal and externally directed attention (Spreng et al., 2013). Chapter 5 found that FPCN connectivity predicted whether a participant was performing a mental imagery task or resting. Chapter 6 used this mental imagery task with DOC patients, but because we were uncertain about their volitional capacity while they were performing the task, we could not reliably label the conditions as mental imagery or resting. We, therefore, used a univariate analysis first to determine which patients were able to perform the task reliably. Next, we used patient resting state connectivity to predict whether they were able to perform the task and found that FPCN connectivity again provided the best accuracy. This is in line with some recent results suggesting that FPCN connectivity is essential in identifying preserved cognitive capacity in DOC patients

(Naci et al., 2014) and further the development of machine learning applications in clinical neuroscience.

Another important aspect of this thesis was to more fully characterise the network neuroscience underlying propofol-induced sedation. To this end, Chapters 3 and 4 examined multi-scale neurobiological questions in healthy subjects undergoing fMRI scanning while receiving an infusion of propofol. In Chapter 3, we found that several large-scale brain networks, most notably the DMN, FPCN, and SN altered their functional relationships during propofol administration. Additionally, we found that this alteration in functional connectivity overlapped with the distribution of parvalbumin-expressing GABAergic interneurons, suggesting they play a role mediating changes in cortical connectivity patterns. In Chapter 4, we specifically examined connectivity within and between several DMN, FPCN, and SN using both standard correlations and graph theoretical metrics. We found that between DMN and SN connectivity increased significantly during sedation. This is significant in that we observed a reduction of anticorrelated edges between network nodes during sedation, a finding also shown in DOC patients (Di Perri et al., 2016). It should be noted that this interactions is different then what we observed in DOC, where FPCN connectivity was significantly different between VC and non-VC patients. There are two reasons why this may be the case; the first is that the neurophysiology associated with loss of consciousness is different for DOC patients and propofol-induced sedation. This fits with integrated information theory in that consciousness is a global phenomenon, and that significant loss of information integration results in loss of consciousness when any network interactions are sufficiently perturbed; the second is that DOC patients were stratified not on their level of consciousness (as reflected in their diagnosis), but by their VC. This may be the reason we see FPCN as more significant, as it is known that this network is important for cognitive control and executive functions. We also found that the structural connections between DMN and SN nodes increased their correspondence significantly with functional connections between these networks. These findings highlight the importance of DMN connectivity, specifically anticorrelations between DMN and attentional systems, during propofol sedation (Boveroux, Vanhaudenhuyse, Bruno, et al., 2010; Stamatakis et al., 2010). The results derived from the gene correlation method from Chapter 3 are

also significant in that they are among the first to use whole-brain microarray data to examine the relationship between functional connectivity and genetic markers for specific neurotransmitter receptors (Preller et al., 2018). Furthermore, the results from Chapter 4 are the first to study structure-function relationships between large-scale brain networks during propofol sedation in humans (Barttfeld et al., 2015). Taken together, these studies expand the current literature on large-scale brain network connectivity during sedation and anaesthesia.

In Chapters 4 and 6 we also observed differences in system segregation between groups. In Chapter 4, we found that sedation reduces network segregation. In Chapter 6 we found that non-VC patients had increased segregation relative to VC patients. This is seemingly contradictory, in that Chapter 4 suggests that less conscious individuals have reduced segregation, while Chapter 6 suggests that less conscious patients have increased segregation. However, this difference may be due to the fact that DOC patients were not stratified based on level of consciousness, but by their ability to perform a volitional mental imagery task in the scanner. Furthermore, when examining the raw system segregation values for each group, we find that VC DOC patients (0.8) were relatively close to awake healthy controls (0.63). This may suggest that there is a critical point where system segregation is optimal.

Finally, this thesis contains several methodological advancements that were developed in order to perform each experiment. Chapter 2 examined the role of temporal filtering in functional connectivity analysis. We found that when examining between network connectivity and task-related functional connectivity that bandpass filtering removes potentially important signal. As task-related functional connectivity in DMN and between network connectivity are becoming increasingly present in the literature, more accurate temporal signal processing is vital to performing optimal analyses. Chapter 3 developed a novel method for identifying relationships between specific neurotransmitter receptors and measures of functional connectivity. Finally, Chapters 5 and 6 used a DGCNN to classify brain graphs into different experimental conditions. Previous work (Ktena et al., 2018; Parisot et al., 2018) has applied DGCNN to brain network data. However, Chapter 5 was the first to apply it using a cognitive paradigm, and Chapter 6 was the first study to apply it to patients with

DOC. Further, as deep learning methods rely on a large number of data points, we applied a novel dynamic connectivity-based data augmentation technique to increase the amount of data available.

7.3. Limitations and Future Directions

As discussed previously, the experimental chapters (Chapters 3 and 4) using propofol to model changes in network connectivity do not use doses high enough to produce unconsciousness. Therefore, the conclusions that can be drawn from this work are limited in that they do not explicitly model complete unconsciousness, but sedation. Previous work suggests that the effect of propofol is graded (Davis et al., 2007), so the present results likely reflect neurobiological processes associated with very low levels of consciousness. Further work could use higher doses of propofol, similar to levels used in general anaesthesia, to induce complete unconsciousness (MacDonald et al., 2015).

In addition to using increased dosages of anaesthetic drugs, data collected using other drugs and other altered states of consciousness could help tease apart state specific, as opposed to drug-specific changes in network connectivity. A recent review calls into question the validity of conceptualising consciousness exclusively as a stratified system (Bayne et al., 2016). They instead develop a multidimensional framework for global states of consciousness, consisting of dimensions for gating conscious content and various functional dimensions related to behaviour and cognition. Therefore, an experimental design more suited to this framework could involve collecting data from the same subjects using different drug states and different levels of arousal (e.g.. sleep). It would also be useful to collect task data in addition to resting state data, as that would allow for the examination of different cognitive and behavioural dimensions within each global state. Data-driven statistical analyses could then be used to identify whether a certain neural pattern changed systematically with the changing of either the global state (the gating dimension) or the cognitive and behavioural state (functional dimension). This type of experiment, though not directly studying DOC patients, would allow for a more fine-grained approach to engineering features as inputs for machine learning approaches to clinical assessment.

Finally, several improvements could be made to the machine learning approaches used in Chapter 6. First, the sample size used here is relatively small (Lecun et al., 2015). This is not surprising, as access to patient data is limited. However, larger multicentre datasets should be used prior to any machine learning use in a clinical setting for several reasons. One, more data will result in a more generalizable model because more extensive hyperparameter searching can be performed and variance between the training, validation, and test sets will be reduced. Second, it is critical to use data collected from different centres, as different scanners can introduce various biases into the dataset (Parisot et al., 2018). In addition to using more data, other imaging modalities including EEG, MEG, and PET should be explored using DGCNNs. In fact, due to DGCNNs having feature vectors at each node, it would be possible to combine data from multiple imaging modalities into a single analysis. With the proper spatial and temporal pre-processing of each imaging modality, one could conceivably concatenate the features from each modality within each node. This would result in a denser feature representation and would likely result in higher classification scores. Another improvement to DGCNN would be more comprehensive thresholding of functional connectivity graphs. The convolutions applied to each node are dependent on its local neighbourhood, making thresholding of the graph particularly important. The present work uses standard correlation cut-offs, however recent work in network neuroscience (Váša, Bullmore, & Patel, 2018) and computational biology (Wang et al., 2018) has developed data-driven methods that could be incorporated as an additional pre-processing step prior to input into the DGCNN. The use of other brain regions as input features would also greatly improve the generalizability of machine learning algorithms to DOC patients. Due to the use of cortical functional connectivity graphs, our analysis was reliant on using patients with relatively limited brain damage. As DOC patients often have significant cortical and subcortical lesions, the use of cortical functional connectivity limits the algorithms capacity to be applied to a wider group of DOC patients. Therefore, it is critical to explore whether the use of other brain regions yields comparable or better results. Recent work using voxel-based lesion-symptom mapping identified a brainstem region, the rostral dorsolateral pontine tegmentum, which was significantly associated with coma. They then used that region as a seed in a functional connectivity analysis

in healthy participants revealing significant connectivity with anterior insula and the anterior cingulate. The authors then found that functional connectivity between these two cortical regions was significantly lower than that of other functional networks, suggesting these connections play an important role in maintaining consciousness (Fischer et al., 2016). Therefore, this set of regions could be an interesting input feature for a DGCNN. Taken together, an improved experiment aiming to differentiate between patients with DOC would include more patients from multiple imaging centres, improved featurisation and pre-processing, and potentially additional imaging modalities. This could then be developed into a clinical tool to aid in the diagnosis and prognosis of patients suffering from DOC.

In sum, a wealth of exciting new findings can be expected from future research aimed at understanding the neurobiology of consciousness. The work presented in this thesis has made a unique contribution to the existing body of knowledge by further characterising network connectivity during propofol sedation and DOC and advancing computational methods to aid in the diagnosis of DOC patients. Finally, and perhaps most importantly, it has set solid grounds and opened up a number of questions for future research to address to better understand the neurobiological basis of consciousness and its importance to clinical neurosciences.

References

Abdi, H. (2010). Partial least squares regression and projection on latent structure regression (PLS Regression). *Wiley Interdisciplinary Reviews: Computational Statistics*, *2*(1), 97–106.

Absalom, A. R., Mani, V., Smet, T. De, & Struys, M. M. (2009). Pharmacokinetic models for propofol — defining and illuminating the devil in the detail, *103*(February), 26–37.

Adapa, R. M., Davis, M. H., Stamatakis, E. A., Absalom, A. R., & Menon, D. K. (2013). Neural Correlates of Successful Semantic Processing During Propofol Sedation. *Human Brain Mapping*, *35*, 2935–2949.

Akeju, O., & Brown, E. N. (2017). Neural oscillations demonstrate that general anesthesia and sedative states are neurophysiologically distinct from sleep. *Current Opinion in Neurobiology*, *44*, 178–185.

Alexander-Bloch, A., Giedd, J. N., & Bullmore, E. (2013). Imaging structural co-variance between human brain regions Aaron. *Nature Reviews Neuroscience*, *14*(5), 322–336.

Alkire, M. T., Hudetz, A. G., & Tononi, G. (2008). Consciousness and Anesthesia, *322*(November), 876–881.

Altman, N. S. (1992). An Introduction to Kernel and Nearest-Neighbor Nonparametric Regression Author(s): *The American Statistician*, *46*(3), 175–185.

Andreasen, N. C., O'Leary, D. S., Cizaldo, T., Arndt, S., Rezai, K., Watkins, G. L., … Hichwa, R. D. (1995). Remembering the past: two facets of episodic memory explored with positron emmision tomogrpahy. *American Journal of Psychiatry*, *152*(11), 1576–1585.

Andrews-Hanna, J. R., Reidler, J. S., Sepulcre, J., Poulin, R., & Buckner, R. L. (2010). Functional-Anatomic Fractionation of the Brain's Default Network. *Neuron*, *65*(4), 550–562.

Andrews-Hanna, J. R., Smallwood, J., & Spreng, R. N. (2014). The default network and self-generated thought : component processes , dynamic control , and clinical relevance. *Annals of the New York Academy of Sciences*, *1316*, 29–52.

Andrews, K., Murphy, L., Munday, R., & Littlewood, C. (1996). Misdiagnosis of the vegetative state : retrospective study in a rehabilitation unit. *British Medical Journal*, *313*(13), 1–8.

Ansell, B. J., & Keenan, J. E. (1989). The Western Neuro Sensory Stimulation Profile : A Tool for Assessing Slow-to-Recover Head-Injured Patients. *Archives of Physical Medicine and Rehabilitation*, *70*, 104–108.

Antognini, J. F., Buonocore, M. H., Disbrow, E. A., & Carstens, E. (1997). Isoflurane Anesthesia Blunts Cerebral Responses to Noxious and Innocuous Stumuli: A fMRI Study. *Pharmacology Letters*, *61*(24), 349–354.

Ashburner, J., & Friston, K. J. (2005). Unified segmentation. *NeuroImage*, *26*, 839–851.

Ashwal, S., Eyman, R. K., & Call, T. L. (1994). Life Expectancy of Children in a Persistent Vegetative State. *Pediatric Neurology*, *10*, 27–33.

Aubinet, C., Larroque, S. K., Heine, L., Martial, C., Majerus, S., Laureys, S., & Di Perri, C. (2018). Clinical subcategorization of minimally conscious state according to resting functional connectivity. *Human Brain Mapping*, *39*, 4519–4532.

Avramov, M. N., White, P. F., & Husain, M. (1995). The Comparative Effects of Methohexital, Etomidate for Electroconvulsive Therapy. *Anesthesia and Analgesia*, *81*, 596–602.

Baars, B. (2002). The conscious access hypothesis : origins and recent evidence. *Trends in Cognitive Sciences*, *6*(1), 47–52.

Bailey, D. L., Townsend, D. W., Valk, P. E., & Maisey, M. N. (2005). *Positron Emission Tomography*.

Bardin, J. C., Schiff, N. D., & Voss, H. U. (2012). Pattern Classification of Volitional Functional Magnetic Resonance Imaging Responses in Patients With Severe Brain Injury. *Archives of Neurology*, *69*(2), 176–181.

Bartos, M., Vida, I., & Jonas, P. (2007). Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature Reviews Neuroscience*, *8* (January), 45–56. https://doi.org/10.1038/nrn2044

Barttfeld, P., Uhrig, L., Sitt, J. D., Sigman, M., Jarraya, B., & Dehaene, S. (2015). Signature of consciousness in the dynamics of resting-state brain activity. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(3), 887–892.

Bassett, D. S., & Sporns, O. (2017). Network neuroscience. *Nature Neuroscience*, *20*(3), 353–364.

Bayne, T., Hohwy, J., & Owen, A. M. (2016). Are There Levels of Consciousness ? *Trends in Cognitive Sciences*, *20*(6), 405–413.

Beaty, R. E., Kenett, Y. N., Christensen, A. P., Rosenberg, M. D., Benedek, M., Chen, Q., … Silvia, P. J. (2018). Robust prediction of individual creative ability from brain functional connectivity. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(5), 1087–1092.

Beaumont, J. G., & Kenealy, P. M. (2005). Incidence and prevalence of the vegetative and minimally conscious states. *Neuropsychological Rehabilitation*, *15*(3/4), 184–189.

Behzadi, Y., Restom, K., Liau, J., & Liu, T. T. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *NeuroImage*, *37*(1), 90–101.

Bekinschtein, T., Manes, F. F., Villarreal, M., Owen, A. M., & Della-Maggiore, V. (2011). Functional imaging reveals movement preparatory activity in the vegetative state. *Frontiers in Human Neuroscience*, *5*, 1–11.

Bell, a J., & Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, *7*(6), 1129–59.

Bergman, S. A. (1999). Ketamine : Review of Its Pharmacology and Its Use in Pediatric Anesthesia Ketamnine. *Anesthesia Progress*, *46*, 10–20.

Bernat, J. L. (2006). Chronic disorders of consciousness. *The Lancet*, *367*, 1181–1192.

Birn, R. M. (2012). The role of physiological noise in resting-state functional connectivity. *NeuroImage*, *62*(2), 864–870.

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*(2), 644–54.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.

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 Functional Connectivity in the Motor Cortex of Resting Human Brain Using Echo-Planar MRI. *Magnetic Resonance Imaging in Medicine*. 34(4), 537-541

Bock, D. D., Allen Lee, W.-C., Kerlin, A. M., Andermann, M. L., Wetzel, A. W., Yurgenson, S., … Reid, C. (2011). Network anatomy and in vivo physiology of visual cortical neurons. *Nature*, *471*(7337), 177–182.

Boly, M., Coleman, M. R., Davis, M. H., Hampshire, A., Bor, D., Moonen, G., … Owen, A. M. (2007). When thoughts become action : An fMRI paradigm to study volitional brain activity in non-communicative brain injured patients. *NeuroImage, 36*, 979–992.

Boly, M., Faymonville, M.-E., Peigneux, P., Lambermont, B., Damas, P., Fiore, G.-D., … Laureys, S. (2004). Auditory Processing in Severely Brain Injured Patients. *Archives of Neurology*, *61*.

Boly, M., Seth, A. K., Wilke, M., Ingmudson, P., Baars, B., Laureys, S., … Tsuchiya, N. (2013). Consciousness in humans and non-human animals : recent advances and future directions. *Frontiers in Psychology*, *4*, 1–20.

Boly, M., Tshibanda, L., Vanhaudenhuyse, A., Noirhomme, Q., Schnakers, C., Ledoux, D., … 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*, 2393–2400.

Bonhomme, V., Fiset, P., Meuret, P., Backman, S., Plourde, G., Paus, T., … Evans, A. C. (2001). Propofol Anesthesia and Cerebral Blood Flow Changes Elicited by Vibrotactile Stimulation : A Positron Emission Tomography Study. *Journal of Neurophysiology*, 1299–1308.

Bonhomme, V., Vanhaudenhuyse, A., Demertzi, A., Bruno, M.-A., Jaquet, O., Bahri, M. A., … Laureys, S. (2016). Functional Connectivity during Ketamine Alteration of Consciousness in Volunteers. *Anesthesiology*, *125*(5), 873–888.

Bonnelle, V., Ham, T. E., Leech, R., Kinnunen, K. M., Mehta, M. A., Greenwood, R. J., & Sharp, D. J. (2012). Salience network integrity predicts default mode network function after traumatic brain injury. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(12), 4690–4695.

Boubela, R. N., Kalcher, K., Hug, W., Kronnerwetter, C., Filzmoser, P., & Moser, E. (2013). Beyond noise : using temporal ICA to extract meaningful information from high-frequency fMRI signal fluctuations during rest. *Frontiers in Human Neuroscience*, *7*(May), 1–12.

Boveroux, P., Vanhaudenhuyse, A., Bruno, M.-A., Noirhomme, Q., Lauwick, S., Luxen, A., … Boly, M. (2010). Breakdown of within- and between-network Resting State during Propofol-induced Loss of Consciousness. *Anesthesiology*, *113*(5), 1038–53.

Boveroux, P., Vanhaudenhuyse, A., & Phillips, C. (2010). Breakdown of within- and between-network Resting State during Propofol-induced Loss of Consciousness. *Anesthesiology*, *113*(5), 1038–1053.

Boyacioglu, R., Beckmann, C. F., & Barth, M. (2013). An investigation of RSN frequency spectra using ultra-fast generalized inverse imaging. *Frontiers in Human Neuroscience*, *7*(April), 1–7.

Braakman, R., Jennett, W. B., & Minderhoud, J. M. (1988). Prognosis of the Posttraumatic Vegetative State. *Acta Neurochirurgica*, *95*, 49–52.

Bressler, S. L., & Menon, V. (2010). Large-scale brain networks in cognition : emerging methods and principles. *Trends in Cognitive Sciences*, *14*(6), 277–290.

Brown, E. N., Lydic, R., & Schiff, N. D. (2010). General Anesthesia, Sleep, and Coma Emery. *New England Journal of Medicine*, *363*(27), 2638–2650.

Brown, E. N., Purdon, P. L., & Van Dort, C. J. (2011). General Anesthesia and Altered States of Arousal: A Systems Neuroscience Analysis. *Annual Review of Neuroscience*, *34*, 601–628.

Bruno, M.-A., Vanhaudenhuyse, A., Thibaut, A., Moonen, G., & Laureys, S. (2011). From unresponsive wakefulness to minimally conscious PLUS and functional locked-in syndromes : recent advances in our understanding of disorders of consciousness. *Journal of Neurology*, *258*, 1373–1384.

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.

Buckner, R. L., & Carroll, D. C. (2007). Self-projection and the brain. *Trends in Cognitive Sciences*, *11*(2), 49–57.

Bullmore, E., & Bassett, D. S. (2011). Brain Graphs : Graphical Models of the Human Brain Connectome. *Annual Review in Clinical Psychology*, *7*, 113–140.

Bullmore, E., & Sporns, O. (2009). Complex brain networks : graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, *10*,

Bunge, S. A., & Kahn, I. (2009). Cognition : An Overview of Neuroimaging Techniques. In *Encylopedia of Neuroscience* (Vol. 2, pp. 1063–1067).

Buzsaki, G. (1983). Cellular bases of hippocampal EGG in behaving rat. *Brain Research*, *173*(April 2016).

Buzsaki, G., & Wang, X. (2012). Mechanisms of Gamma Oscillations. *Annual Review of Neuroscience*, *35*, 203–225.

Bzdok, D., Heeger, A., Langner, R., Laird, A. R., Fox, P. T., Palomero-gallagher, N., … Eickhoff, S. B. (2015). Subspecialization in the human posterior medial cortex. *NeuroImage*, *106*, 55–71.

Carollo, D. S., Nossaman, B. D., & Ramadhyani, U. (2008). Dexmedetomidine : a review of clinical applications. *Current Opinion in Anesthesiology*, *21*, 457–461.

Carr, M. F., Karlsson, M. P., & Frank, L. M. (2012). Transient Slow Gamma Synchrony Underlies Hippocampal Memory Replay. *Neuron*, 75(4), 700-713.

Cauda, F., Micon, B. M., Sacco, K., Duca, S., Agata, F. D., Geminiani, G., & Canavero, S. (2009). Disrupted intrinsic functional connectivity in the vegetative state. *Journal of Neurology, Neurosurgery and Psychiatry*, *80*, 429–431.

Chai, X. J., Nieto-Castañón, A., Ongür, D., & Whitfield-Gabrieli, S. (2012). Anticorrelations in resting state networks without global signal regression. *NeuroImage*, *59*(2), 1420–8.

Chalmers, D. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200-219

Chan, M. T., Park, D. C., Savalia, N. K., Petersen, S. E., & Wig, G. S. (2014). Decreased segregation of brain systems across the healthy adult lifespan. *Proceedings of the National Academy of Sciences of the United States of America*, E4997–E5006.

Chang, C.-C., & Lin, C.-J. (2011). LIBSVM : A Library for Support Vector Machines. In *ACM Transactions on Intelligent Systems and Technology* (pp. 1–39).

Chen, C.-H., Panizzon, M. S., Eyler, L. T., Jernigan, T. L., Fennema-Notestine, C., Jak, A. J., … Larry, J. (2011). Genetic Influences on Cortical Regionalization in the Human Brain. *Neuron*, *72*(4), 537–544.

Chen, J. E., & Glover, G. H. (2015). BOLD fractional contribution to resting-state functional connectivity above 0.1 Hz. *NeuroImage*, *107*, 207–218.

Chen, Z. J., He, Y., Rosa-Neto, P., Germann, J., & Evans, A. C. (2008). Revealing Modular Architecture of Human Brain Structural Networks by Using Cortical Thickness from MRI. *Cerebral Cortex*, *18*, 2374–2381.

Chennu, S., Annen, J., Wannex, S., Thibaut, A., Chatelle, C., Casol, H., … Laureys, S. (2017). Brain networks predict metabolism, diagnosis and prognosis at the bedside in disorders of consciousness. *Brain*, *140*, 2120–2132.

Chennu, S., Finoia, P., Kamau, E., Allanson, J., Williams, G. B., Monti, M. M., … Arnatkeviciute, A. (2014). Spectral Signatures of Reorganised Brain Networks in Disorders of Consciousness. *PLoS Computational Biology*, (10), e1003887.

Chennu, S., O'Connor, S., Adapa, R. M., Menon, D. K., & Bekinschtein, T. (2016). Brain Connectivity Dissociates Responsiveness from Drug Exposure during Propofol-Induced Transitions of Consciousness. *PLoS Computational Biology*, *12*(1), 1–17.

Chennu, S., Stamatakis, E. A., & Menon, D. K. (2016). The see-saw brain : recovering consciousness after brain injury. *The Lancet Neurology*, *15*(8), 781–782.

Childs, N. L., Mercer, W. N., & Childs, H. W. (1993). Accuracy of diagnosis of persistent vegetative state. *Neurology*, *43*, 1465–1467.

Chiu, T.-H., Chen, M.-J., Yang, Y.-R., Yang, J.-J., & Tang, F. –. (1995). Action of dexmedetomidine on rat locus coeruleus neurones: intracellular recording in vitro. *European Journal of Pharmacology*, *285*, 261–268.

Choi, S. C., Barnes, T. Y., Bullock, R., Germanson, T. A., Marmarou, A., & Young, H. F. (1994). Temporal profile of outcomes in severe head injury. *Jounal of Neurosurgery*, *81*, 169–173.

Cohen-Adad, J., Descoteaux, M., Rossignol, S., Hoge, R. D., Deriche, R., & Benali, H. (2008). Detection of multiple pathways in the spinal cord using q-ball imaging. *NeuroImage*, *42*, 739–749.

Cohen, J. R., & D'Esposito, M. (2016). The Segregation and Integration of Distinct Brain Networks and Their Relationship to Cognition. *The Journal of Neuroscience*, *36*(48), 12083–12094.

Cole, M. W., Bassett, D. S., Power, J. D., Braver, T. S., & Petersen, S. E. (2014). Intrinsic and Task-Evoked Network Architectures of the Human Brain. *Neuron*, *83*(1), 238–251.

Cole, M. W., Ito, T., Bassett, D. S., & Schultz, D. H. (2016). Activity flow over resting-state networks shapes cognitive task activations. *Nature Neuroscience*, *19*(12), 1718–1729.

Coleman, M. R., Rodd, J. M., Davis, M. H., Johnsrude, I. S., Menon, D. K., Pickard, J. D., & Owen, A. M. (2007). Do vegetative patients retain aspects of language comprehension? Evidence from fMRI. *Brain*, *130*(10), 2494–2507.

Correia, M. M., Carpenter, T. A., & Williams, G. B. (2009). Looking for the optimal DTI acquisition scheme given a maximum scan time : are more b-values a waste

of time ? *Magnetic Resonance Imaging*, *27*(2), 163–175.

Coursin, D. B., Coursin, D. B., & Maccioli, G. A. (2001). Dexmedetomidine. *Current Opinion in Critical Care*, *7*, 221–226.

Crick, F., & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in The Neurosciences*, *2*, 263–275.

Crittenden, B. M., Mitchell, D. J., & Duncan, J. (2015). Recruitment of the default mode network during a demanding act of executive control. *eLife*, *4*, e06481.

Cruse, D., Chennu, S., Chatelle, C., Bekinschtein, T., Fernández-espejo, D., Pickard, J. D., … Owen, A. M. (2011). Bedside detection of awareness in the vegetative state : a cohort study. *The Lancet*, *378*(9809), 2088–2094.

Currie, S., Hoggard, N., Craven, I. J., Hadjivassiliou, M., & Wilkinson, I. D. (2013). Understanding MRI : basic MR physics for physicians. *Postgraduate Medical Journal*, *89*, 209–223.

Dahan, A., Aarts, L., & Smith, T. (2010). Incidence, Reversal, and Prevention of Opioid-induced Respiratory Depression. *Anesthesiology*, *112*, 226–238.

Davey, C. E., Grayden, D. B., Egan, G. F., & Johnston, L. (2013). Filtering induces correlation in fMRI resting state data. *NeuroImage*, *64*(2013), 728–40.

Davis, J., & Goadrich, M. (2006). The Relationship Between Precision-Recall and ROC Curves. In *Proceedings of the 23rd International Conference on Machine Learning*.

Davis, M. H., Coleman, M. R., Absalom, A. R., Rodd, J. M., Johnsrude, I. S., Matta, B. F., … Menon, D. K. (2007). Dissociating speech perception and comprehension at reduced levels of awareness. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(41), 16032–16037.

De Domenico, M., Sasai, S., & Arenas, A. (2016). Mapping Multiplex Hubs in Human Functional Brain Networks. *Frontiers in Neuroscience*, *10*(July), 1–14.

de Jong, B. M., Willemsen, A. T. M., & Paans, A. M. J. (1997). Regional cerebral blood flow changes related to affective speech presentation in persistent vegetative state. *Clinical Neurology and Neurosurgery*, *99*, 213–216.

Deco, G., Jursa, V., McIntosh, A. R., Sporns, O., & Kotter, R. (2009). Key role of coupling, delay, and noise in resting brain fluctuations. *Proceedings of the National Academy of Sciences of the United States of America*, 106(29),10302–10307

Deco, G., Kringelbach, M. L., Jirsa, V. K., & Ritter, P. (2017). The dynamics of resting fluctuations in the brain: metastability and its dynamical cortical core. *Scientific Reports*, *7*(April), 1–14.

Deco, G., Tononi, G., Boly, M., & Kringelbach, M. L. (2015). Rethinking segregation and integration: contributions of whole-brain modelling. *Nature Publishing Group*, *16*(7), 430–439.

Dehaene, S., Changeux, J.-P., & Naccache, L. (2011). The Global Neuronal Workspace Model of Conscious Access: From Neuronal Architectures to Clinical Applications. In S. Dehaene & Y. Christen (Eds.), *Characterizing Consciousness: From Cognition to the Clinic? Research* (pp. 55–84). Berlin, Heidelberg: Springer Berlin Heidelberg.

Dehaene, S., Charles, L., King, J.-R., & Marti, S. (2014). Toward a computational theory of conscious processing. *Current Opinion in Neurobiology*, *25*, 76–84.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness : basic evidence and a workspace framework. *Cognition*, *79*, 1–37.

Della-Maggiore, V., Chau, W., Peres-Neto, P. R., & McIntosh, A. R. (2002). An Empirical Comparison of SPM Preprocessing Parameters to the Analysis of fMRI Data. *NeuroImage*, *17*(1), 19–28.

Demertzi, A., Antonopoulos, G., Heine, L., Voss, H. U., Crone, J. S., de Los Angeles, C., … Laureys, S. (2015). Intrinsic functional connectivity differentiates minimally conscious from unresponsive patients. *Brain*, *138*, 2619–2631.

Demertzi, A., Gomez, F., Crone, J. S., Vanhaudenhuyse, A., Tshibanda, L., Noirhomme, Q., … Soddu, A. (2014). Multiple fMRI system-level baseline connectivity is disrupted in patients with consciousness alterations. *Cortex*, *2*, 35–46.

Demertzi, A., Soddu, A., & Laureys, S. (2013). Consciousness supporting networks. *Current Opinion in Neurobiology*, *23*(2), 239–44.

Di Perri, C., Bahri, M. A., Amico, E., Thibaut, A., Heine, L., Antonopoulos, G., … Laureys, S. (2016). Neural correlates of consciousness in patients who have emerged from a minimally conscious state: a cross-sectional multimodal imaging study. *The Lancet Neurology*, *15*(8), 830–42.

Dixon, M. L., Andrews-Hanna, J. R., Spreng, R. N., Irving, Z. C., Mills, C., Girn, M., & Christoff, K. (2017). NeuroImage Interactions between the default network and dorsal attention network vary across default subsystems , time , and cognitive

states. *NeuroImage*, *147*(December), 632–649.

Dosenbach, N. U. F., Fair, D. A., Miezin, F. M., Cohen, A. L., Wenger, K. K., Dosenbach, R. A. T., … Petersen, S. E. (2007). Distinct brain networks for adaptive and stable task control in humans. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(26), 11073–11078.

Dosenbach, N. U. F., Visscher, K. M., Palmer, E. D., Miezin, F. M., Wenger, K. K., Kang, H. C., … Petersen, S. E. (2006). A Core System for the Implementation of Task Sets. *Neuron*, *50*, 799–812.

Dubroja, I., Valent, S., Miklic, P., & Kesak, D. (1995). Outcome of post-traumatic unawareness persisting for more than a month. *Journal of Neurology, Neurosurgery and Psychiatry*, *58*, 465–466.

Eckart, C., & Young, G. (1936). The approximation of one matrix by another of lower rank. *Psychometrika*, *1*(3), 211–218.

Efron, B., & Tibshirani, R. (1986). Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy. *Statistical Science*, *1*(1), 54–77.

Feinberg, D. A., Moeller, S., Smith, S. M., Auerbach, E., Ramanna, S., Matt, F., … Yacoub, E. (2010). Multiplexed Echo Planar Imaging for Sub-Second Whole Brain FMRI and Fast Diffusion Imaging. *PloS One*, *5*(12), e15710.

Fernández-espejo, D., Bekinschtein, T., Monti, M. M., Pickard, J. D., Junque, C., Coleman, M. R., & Owen, A. M. (2011). Diffusion weighted imaging distinguishes the vegetative state from the minimally conscious state. *NeuroImage*, *54*(1), 103–112.

Fernández-Espejo, D., Norton, L., & Owen, A. M. (2014). The Clinical Utility of fMRI for Identifying Covert Awareness in the Vegetative State: A Comparison of Sensitivity between 3T and 1.5T. *PLOS ONE*, *9*(4), e95082.

Fins, J. J., & Shapiro, Z. E. (2007). Neuroimaging and neuroethics : clinical and policy considerations. *Current Opinion in Neurology*, *20*, 650–654.

Fischer, D. B., Boes, A. D., Demertzi, A., Evrard, H. C., Laureys, S., Edlow, B. L., … Geerling, J. C. (2016). A human brain network derived from coma-causing brainstem lesions. *Neurology*, *87*, 2427–2434.

Fox, K. C. R., Spreng, R. N., Ellamil, M., Andrews-Hanna, J. R., & Christoff, K. (2015). The wandering brain: Meta-analysis of functional neuroimaging studies of mind-wandering and related spontaneous thought processes. *NeuroImage*, 1(11), 611-

621.

Fox, M. D., & Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging, *Nature Reviews Neuroscience*, 8(9), 700-711.

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*(27), 9673–9678.

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*, *29*, 15–29.

Friston, K. J. (1994). Functional and Effective Connectivity in Neuroimaging : A Synthesis. *Human Brain Mapping*, 2, 56-78.

Galanaud, D., Perlbarg, V., Gupta, R., Stevens, R. D., Sanchez, P., Tollard, E., … Puybasset, L. (2012). Assessment of White Matter Injury and Outcome in Severe Brain Trauma. *Anesthesiology, 2(6)*, 1300–1310.

Gerlach, A. T., & Dasta, J. F. (2007). Dexmedetomidine : An Updated Review. *Critical Care*, *41*, 245–253.

Gerlach, K. D., Spreng, R. N., Madore, K. P., & Schacter, D. L. (2014). Future planning : default network activity couples with frontoparietal control network and reward-processing regions during process and outcome simulations. *Social, Cognitive, and Affective Neuroscience, 9(12)*, 1942–1951.

Giacino, J. T., Ashwal, S., Childs, N., Cranford, R., Jennett, B., Katz, D. I., … Zasler, N. D. (2002). The minimally conscious state: Definition and diagnostic criteria. *Neurology*, *58*, 349–354.

Giacino, J. T., Fins, J. J., Laureys, S., & Schiff, N. D. (2014). Disorders of consciousness after acquired brain injury: the state of the science. *Nature Reviews Neurology*, *10*(2), 99–114.

Giacino, J. T., & Kalmar, K. (2011). Diagnostic and prognostic guidelines for the vegetative and minimally conscious states. *Neuropsychological Rehabilitation*, *15*, 166–174.

Giacino, J. T., Kalmar, K., & Whyte, J. (2004). The JFK Coma Recovery Scale – Revised : Measurement. *Archives of Physical Medicine and Rehabilitation*, *85*, 2020–2029.

Gill-Thwaites, H., & Munday, R. (2011). The Sensory Modality Assessment and Rehabilitation Technique (SMART): A Comprehensive and Integrated Assessment and Treatment Protocol for the Vegetative State and Minimally Responsive Patient The Sensory Modality Assessment and Rehabilitation Technique. *Neuropsychological Rehabilitation*, *9*, 305–320.

Glasser, M. F., Smith, S. M., Marcus, D. S., Andersson, J. L. R., Auerbach, E. J., Behrens, T. E. J., … Van Essen, D. C. (2016). The Human Connectome Project ' s neuroimaging approach. *Nature Neuroscience*, *9*, 1175–1187.

Gohel, S. R., & Biswal, B. B. (2015). Functional Integration Between Brain Regions. *Brain Connectivity*, *5*(1), 23–34.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.

Greicius, M. D., Kiviniemi, V., Tervonen, O., Vainionpää, V., Reiss, A. L., & Menon, V. (2008). Persistent Default-Mode Network Connectivity During Light Sedation. *Human Brain Mapping*, *29*(7), 839–847.

Greicius, M. D., Krasnow, B., Boyett-Anderson, J. M., Eliez, S., Schatzberg, A. F., Reiss, A. L., & Menon, V. (2003). Regional analysis of hippocampal activation during memory encoding and retrieval: fMRI study. *Hippocampus*, *13*(1), 164–74.

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*, 253–258.

Guimera, R., & Nunes Amaral, L. A. (2005). Functional cartography of complex metabolic networks. *Nature*, *433*(7028), 895–900.

Gusnard, D. A., & Raichle, M. E. (2001). Searching for a baseline: Functional imaging and the resting human brain. *Nature Reviews Neuroscience*, *2*(10), 685–694.

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., & Sporns, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS Biology*, *6*(7), e159.

Han, H. J., Kim, E. J., Lee, H. J., Pyun, S. B., Joa, K. L., & Jung, H. Y. (2018). Validation of Korean Version of Coma Recovery Scale-Revised (K-CRSR). *Annals of Rehabilitation Medicine*, *42*(4), 536–541.

Hannawi, Y., Lindquist, M. A., Caffo, B. S., Sair, H. I., & Stevens, R. D. (2015). Resting brain activity in disorders of consciousness: A systematic review and meta-

analysis. *Neurology*, *84*, 1272–1280.

Hashmi, J. A., Loggia, M. L., Khan, S., Gao, L., Kim, J., Napadow, V., … Akeju, O. (2017). Dexmedetomidine Disrupts the Local and Global Efficiencies of Large-scale Brain Networks. *Anesthesiology*, *126*, 419–430.

Haupt, W. F., & Rudolf, J. (1999). European brain death codes: a comparison of national guidelines. *Journal of Neurology*, *246*, 432–437.

Hawrylycz, M. J., Lein, E. S., Guillozet-Bongaarts, A. L., Shen, E. H., Ng, L., Miller, J. A., … Jones, A. R. (2012). An anatomically comprehensive atlas of the adult human brain transcriptome. *Nature*, *489*, 391–399.

Heinke, W., & Koelsch, S. (2005). The effects of anesthetics on brain activity and cognitive function. *Current Opinion in Anesthesiology*, *18*, 625–631.

Hofbauer, R. K., Fiset, P., Plourde, G., & Backman, S. B. (2004). Dose-dependent Effects of Propofol on the Central Processing of Thermal Pain. *Anesthesiology*, *100*(2), 386–394.

Honey, C. J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J. P., Meuli, R., & Hagmann, P. (2009). Predicting human resting-state functional connectivity. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(6), 2035–2040.

Hsu, Y. W., Cortinez, L. I., Robertson, K. M., J.C., K., Sum-Ping, S. T., Moretti, E. W., … Somma, J. (2004). Dexmedetomidine Pharmacodynamics : Part I Crossover Comparison of the Respiratory Effects of Dexmedetomidine and. *Anesthesiology*, *101*(5), 1066–1076.

Huang, Y., Yoon, K., Ko, H., Jiao, S., Ito, W., Wu, J.-Y., … Morozov, A. (2016). 5-HT3a Receptors Modulate Hippocampal Gamma Oscillations by Regulating Synchrony of Parvalbumin-Positive Interneurons. *Cerebral Cortex*, *26*(2), 576–585.

Hudetz, A. G. (2012). General Anesthesia and Human Brain Connectivity. *Brain Connectivity*, *2*(6), 291–302.

Ito, H., Nyberg, S., Halldin, C., Lundkvist, C., & Farde, L. (1998). PET Imaging of Central 5-HT2A Receptors with. *The Journal of Nuclear Medicine*, *39*(1), 208–214.

Jacobs, J., Weidemann, C. T., Miller, J. F., Solway, A., Burke, J. F., Wei, X.-X., … Kahana, M. J. (2013). Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature Neuroscience*, (A), 8–11.

https://doi.org/10.1038/nn.3466

Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). FSL. *NeuroImage*, *62*, 782–790.

Jiang, X., Shen, S., Cadwell, C. R., Berens, P., Sinz, F., Ecker, A. S., … Tolias, A. S. (2015). Principles of connectivity among morphologically defined cell types in adult neocortex. *Science*, *350*(6264), aac9462.

John, E. R., Prichep, L. S., Aubert, E., & Gugino, L. D. (2001). Invariant Reversible QEEG Effects of Anesthetics. *Consciousness and Cognition, 183*, 165–183.

Jurd, R., Arras, M., Lambert, S., Drexler, B., Siegwart, R., Crestani, F., … Rudolph, U. (2002). General anesthetic actions in vivo strongly attenuated by a point mutation in the GABA(A) receptor beta3 subunit. *The FASEB Journal, 17(2),* 250-252.

Kahn, I., Knoblich, U., Desai, M., Bernstein, J., Graybiel, A. M., Boyden, E. S., … Moore, C. I. (2013). Optogenetic drive of neocortical pyramidal neurons generates fMRI signals that are correlated with spiking activity. *Brain Research*, *1511*, 33–45.

Kaisti, K. K., Oikonen, V., Aalto, S., Hinkka, S., & Scheinin, H. (2002). Effects of Surgical Levels of Propofol and Sevoflurane Anesthesia on Cerebral Blood Flow in Healthy Subjects Studied with Positron Emission Tomography. *Anesthesiology*, *96*(6), 1358–1370.

Kalcher, K., Boubela, R. N., Huf, W., Bartova, L., Kronnerwetter, C., Derntl, B., … Moser, E. (2014). The Spectral Diversity of Resting-State Fluctuations in the Human Brain. *PLoS ONE*, *9*(4), e93375.

Kampfl, A., Schmutzhard, E., Franz, G., Pfausler, B., Haring, H.-P., Ulmer, H., … Aichner, F. (1998). Prediction of recovery from post-traumatic vegetative state with cerebral magnetic-resonance imaging. *The Lancet*, *351*, 1763–1767.

Kawaguchi, Y., & Kubota, Y. (1997). GABAergic Cell Subtypes and their Synaptic Connections in Rat Frontal Cortex. *Cerebral Cortex*, *7*(6), 476–486.

Kay, K. N., Rokem, A., Winawer, J., Dougherty, R. F., Wandell, B. A., & Gonzalez-Castillo, J. (2013). GLMdenoise : a fast , automated technique for denoising task-based fMRI data. *Frontiers in Neuroscience*, *7*(12), 1–15.

Kerssens, C., Hamann, S., Peltier, S., Hu, X. P., Byas-Smith, M. G., & Sebel, P. S. (2005). Attenuated Brain Response to Auditory Word Stimulation with Sevoflurane. *Anesthesiology*, *103*(1), 11–19.

King, J.-R., Sitt, J. D., Faugeras, F., Rohaut, B., El Karoui, I., Cohen, L., … Dehaene, S. (2013). Report Information Sharing in the Brain Indexes Consciousness in Noncommunicative Patients. *Current Biology*, *23*, 1914–1919.

Kirchberger, K., Schmitt, H., Hummel, C., Peinemann, A., Pauli, E., Kettenmann, B., & Stefan, H. (1998). Clonidine- and Methohexital-Induced Epileptiform Discharges Detected by Magnetoencephalography (MEG) in Patients with Localization-Related Epilepsies. *Epilepsia*, *39*(10), 1104–1112.

Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: progress and problems. *Nature Reviews Neuroscience*, *17(5)*, 307–321.

Kotchoubey, B., Kaiser, J., Bostanov, V., & Lutzenberger, W. (2009). Recognition of affective prosody in brain-damaged patients and healthy controls : A neurophysiological study using EEG and whole-head MEG. *Cognitive, Affective and Behavioural Neuroscience*, *9*(2), 153–167.

Krishnan, A., Williams, L. J., Randal, A., & Abdi, H. (2011). Partial Least Squares (PLS) methods for neuroimaging : A tutorial and review. *NeuroImage*, *56*(2), 455–475.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. In *Conference on Neural Information Processing Systems* (1–9).

Ktena, S. I., Parisot, S., Ferrante, E., Rajchl, M., Lee, M., Glocker, B., & Rueckert, D. (2018). Metric learning with spectral graph convolutions on brain connectivity networks. *NeuroImage*, *169*(12), 431–442.

Kubota, Y. (2014). Untangling GABAergic wiring in the cortical microcircuit. *Current Opinion in Neurobiology*, *26*, 7–14.

Landau, W. M., Freygang, W. H., Roland, L. P., Sokoloff, L., & Kety, S. (1955). The local circulation of the living brain: Values in the unanesthetized and anesthetized cat. *Transactions of the American Neurological Association*, 125–129.

Laureys, S., Boly, M., & Maquet, P. (2006). Tracking the recovery of consciousness from coma. *The Journal of Clinical Investigation*, *116*(7), 1823–1825.

Laureys, S., Celesia, G. G., Cohadon, F., Lavrijsen, J. C. M., León-Carrión, J., Sannita, W. G., … Dolce, G. (2010). Unresponsive wakefulness syndrome : a new name for the vegetative state or apallic syndrome. *BMC Medicine*, *8*(1), 1–4.

Laureys, S., Faymonville, M.-E., Degueldre, C., Fiore, G.-D., Damas, P., Lambermont,

B., … Maquet, P. (2000). Auditory processing in the vegetative state. *Brain*, *123*, 1589–1601.

Laureys, S., Faymonville, M. E., Peigneux, P., Damas, P., Lambermont, B., Fiore, G.-D., … Maquet, P. (2002). Cortical Processing of Noxious Somatosensory Stimuli in the Persistent Vegetative State. *NeuroImage*, *17*, 732–741.

Laureys, S., Goldman, S., Phillips, C., Van Bogaert, P., Aerts, J., Luxen, A., … Maquet, P. (1999). Impaired Effective Cortical Connectivity in Vegetative State : Preliminary Investigation Using PET. *NeuroImage*, *382*, 377–382.

Laureys, S., Owen, A. M., & Schiff, N. D. (2004). Review Brain function in coma , vegetative state, and related disorders. *The Lancet Neurology*, *3*, 537–546.

Laureys, S., Perrin, F., & Bredart, S. (2007). Self-consciousness in non-communicative patients. *Consciousness and Cognition*, *16*, 722–741.

Laureys, S., Perrin, F., Schnakers, C., Boly, M., & Majerus, S. (2005). Residual cognitive function in comatose , vegetative and minimally conscious states. *Current Opinion in Neurobiology*, *18*, 726–733.

Laureys, S., & Schiff, N. D. (2012). Coma and consciousness: Paradigms (re)framed by neuroimaging. *NeuroImage*, *61*(2), 478–491.

Lavrijsen, J. C. M., van den Bosch, J. S. G., Koopmans, R. T. C. M., & van Weel, C. (2005). Prevalence and characteristics of patients in a vegetative state in Dutch nursing homes. *Journal of Neurology, Neurosurgery and Psychiatry*, *76*, 1420–1424.

Lecun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*, 436–444.

Lee, J. H., Durand, R., Gradinaru, V., Zhang, F., Goshen, I., Kim, D.-S., … Deisseroth, K. (2010). Global and local fMRI signals driven by neurons defined optogenetically by type and wiring. *Nature*, *465*(7299), 788–792.

Lee, U., Arbor, A., Ku, S., Noh, G., Baek, S., Choi, B., … Arbor, A. (2013). Disruption of Frontal-Parietal Communication by Ketamine, Propofol, and Sevoflurane. *Anesthesiology*, *118*(6), 1264–1275.

Leech, R., Kamourieh, S., Beckmann, C. F., & Sharp, D. J. (2011). Fractionating the default mode network: Distinct contributions of the ventral and dorsal posterior cingulate cortex to cognitive control. *Journal of Neuroscience*, *31*(9), 3217–3224.

Leech, R., & Sharp, D. J. (2014). The role of the posterior cingulate cortex in cognition and disease. *Brain : A Journal of Neurology*, *137*(Pt 1), 12–32.

Lerch, J. P., van der Kouwe, A. J. W., Raznahan, A., Paus, T., Johansen-Berg, H., Miller, K. L., … Sotiropoulos, S. N. (2017). Studying neuroanatomy using MRI, *20*(3).

Lewis, L. D., Setsompop, K., Rosen, B. R., & Polimeni, J. R. (2016). Fast fMRI can detect oscillatory neural activity in humans. *Proceedings of the National Academy of Sciences of the United States of America*. *113(43)*, E6687-E6685.

Lewis, L. D., Weiner, V. S., Mukamel, E. A., Donoghue, J. A., & Eskandar, E. N. (2012). Rapid fragmentation of neuronal networks at the onset of propofol-induced unconsciousness, *109*(49), E3377-E3386.

Liu, X., Lauer, K. K., Ward, B. D., Rao, S. M., Li, S.-J., & Hudetz, A. G. (2013). Propofol Disrupts Functional Interactions between Sensory and High-Order Processing of Auditory Verbal Memory. *Human Brain Mapping*, *33*(10), 2487–2498.

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, *453*(7197), 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*, 150–157.

Lombardi, F., Taricco, M., De Tanti, A., Telaro, E., & Liberati, A. (2002). Sensory stimulation of brain-injured individuals in coma or vegetative state : results of a Cochrane systematic review. *Clinical Rehabilitation*, *16*, 464–472.

Luyt, C.-E., Galanaud, D., Perlbar, V., Vanhaudenhuyse, A., Stevens, R. D., Gupta, R., … Puybasset, L. (2012). Diffusion Tensor Imaging to Predict Long-term Outcome after Cardiac Arrest. *Anesthesiology*, *6*, 1311–1321.

MacDonald, A. A., Naci, L., MacDonald, P. A., & Owen, A. M. (2015). Anesthesia and neuroimaging : investigating the neural correlates of unconsciousness. *Trends in Cognitive Sciences*, *19*(2), 100–107.

Majerus, S., Bruno, M.-A., Schnakers, C., & Giacino, J. T. (2009). The problem of aphasia in the assessment of consciousness in brain damaged patients. *Progress in Brain Research*, *177*, 49–61.

Majerus, S., Vanr De Linden, M., & Shiel, A. (2000). Wessex Head Injury Matrix and Glasgow/Glasgow- Liege Coma Scale: A Validation and Comparison Study. *Neuropsychological Rehabilitation*, *10*, 167–184.

Marsh, B., White, M., Morton, N., & Kenny, G. N. C. (1991). Pharmacokinetic Model

Driven Infusino of Propofol in Children. *British Journal of Anaesthesia*, *67*(1), 41–48.

Martuzzi, R., Ramachandran, R., Qiu, M., Rajeevan, N., & Constable, T. (2011). Functional connectivity and alterations in baseline brain state in humans. *NeuroImage*, *49*(1), 823–834.

Mashour, G. A. (2004). Consciousness unbound: toward a paradigm of general anesthesia. *Anesthesiology*, *100*(2), 428–433.

Mashour, G. A., & Hudetz, A. G. (2018). Neural Correlates of Unconsciousness in Large-Scale Brain Networks. *Trends in Neurosciences*, *41*(3), 150–160.

Mazoyer, B., Zago, L., Mellet, E., Bricogne, S., Etard, O., Houde, O., … Tzourio-Mazoyer, N. (2001). Cortical networks for working memory and executive functions sustain the conscious resting state in man. *Brain Research Bulletin*, *54*(3), 287–298.

McIntosh, A. R. (1999). Mapping Cognition to the Brain Through Neural Interactions. *Memory*, *7*, 523–548.

McIntosh, A. R., & Lobaugh, N. J. (2004). Partial least squares analysis of neuroimaging data: applications and advances. *NeuroImage*, *23*, 250-263.

McIntosh, A. R., & Mišić, B. (2013). Multivariate statistical analyses for neuroimaging data. *Annual Review of Psychology*, *64*, 499–525.

McRobbie, D. W., & Moore, E. A. (2003). *MRI from Picture to Proton*. Cambridge University Press

Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Cardin, J. A., Carle, M., … Tsai, L. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses, *459*(6), 663–668.

Menon, D. K., Owen, A. M., Williams, E. J., Minhas, P. S., Allen, C. M. S., Boniface, S. J., & Pickard, J. D. (1998). Cortical processing in persistent vegetative state. *The Lancet*, *352*, 200.

Metriyakool, K. (1981). Seizures induced by methohexital. *Anesthesiology*, *55*, 718.

Mhuircheartaigh, R. N., Warnaby, C., Rogers, R., Jbabdi, S., & Tracey, I. (2013). Slow-Wave Activity Saturation and Thalamocortical Isolation During Propofol Anesthesia in Humans, *5*(208), 1–9.

Millan, M. J. (2002). Descending control of pain. *Progress in Neurobiology*, *66*, 355–

474.

Miller, G. (2005). What Is the Biological Basis of Consciousness. *Science*, *309*, 79.

Mišić, B., Betzel, R. F., de Reus, M. A., van den Heuvel, M. P., Berman, M. G., McIntosh, A. R., & Sporns, O. (2016). Network-Level Structure-Function Relationships in Human Neocortex. *Cerebral Cortex*, *26(7)*, 3285-3296.

Mišić, B., & Sporns, O. (2017). From Regions to Connections and Networks: New Bridges between Brain and Behavior. *Current Opinion in Neurobiology*, *40*, 1–7.

Monti, M. M., Pickard, J. D., & Owen, A. M. (2012). Visual Cognition in Disorders of Consciousness : from V1 to Top-Down Attention. *Human Brain Mapping*, *34(6),* 1245-1253.

Monti, M. M., Vanhaudenhuyse, A., Coleman, M. R., Boly, M., Pickard, J. D., Tshibanda, L., … Laureys, S. (2010). Willful Modulation of Brain Activity in Disorders of Consciousness. *New England Journal of Medicine*, *362*(7), 579–589.

Morgan, C. J. A., Muetzelfeldt, L., & Curran, H. V. (2009). Ketamine use , cognition and psychological wellbeing : a comparison of frequent , infrequent and ex-users with polydrug and non-using controls. *Addiction*, *104*, 77–87.

Mullie, A., Verstringe, P., Buylaert, W., Houbrechts, H., Michema, N., Delooz, H., … Lewi, P. (1988). Predictive Value of Glasgow Coma Score For Awakening After Out-of-Hospital Cardiac Arrest. *The Lancet*, 23(1), 137–140.

Murphy, K., Birn, R. M., & Bandettini, P. A. (2013). Resting-state fMRI confounds and cleanup. *NeuroImage*, *80*, 349–59.

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? Kevin. *NeuroImage*, *44*(3), 893–905.

Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*.

Naci, L., Cusack, R., Anello, M., & Owen, A. M. (2014). A common neural code for similar conscious experiences in different individuals. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(39), 14277–14282.

Naci, L., Graham, M., Owen, A. M., & Weijer, C. (2016). Covert narrative capacity : Mental life in patients thought to lack consciousness. *Annals of Clinical and Translational Neurology*, *4*(1), 61–70.

Newberg, A., Alavi, A., & Reivich, M. (2002). Determination of Regional Cerebral

Function With FDG-PET Imaging in Neuropsychiatric Disorders. *Seminars in Nuclear Medicine, 32(1)*, 13-34.

Niessing, J., Ebisch, B., Schmidt, K. E., Niessing, M., Singer, W., & Galuske, R. A. W. (2005). Hemodynamic Signals Correlate Tightly with Synchronized Gamma Oscillations. *Science*, *309(5736)*, 948–952.

Noirhomme, Q., Brecheisen, R., Lesenfants, D., Antonopoulos, G., & Laureys, S. (2017). "Look at my classifier's result": Disentangling unresponsive from (minimally) conscious patients Quentin. *NeuroImage*, *145*, 288–303.

O'Dell, M. W., Jasin, P., Lyons, N., Stivers, M., & Meszaros, F. (1996). Standardized assessment instruments for minimally-responsive, brain-injured patients. *Neurorehabilitation*, *6*, 45–55.

Ogawa, S., Lee, T. M., Kay, A. R., & Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences of the United States of America*, *87*, 9868–9872.

Ogawa, S., Menon, R. S., Kim, S.-G., & Ugurbil, K. (1998). On the characteristics of functional magnetic resonance imaging of the brain. *Annual Reivew of Biophysics*, *27*, 447–474.

Olney, J. W., Newcomer, J. W., & Farber, N. B. (1999). NMDA receptor hypofunction model of schizophrenia. *Journal of Psychiatric Research*, *33*, 523–533.

Ovadia-Caro, S., Nir, Y., Soddu, A., Ramot, M., Hesselmann, G., Vanhaudenhuyse, A., … Malach, R. (2012). Reduction in Inter-Hemispheric Connectivity in Disorders of Consciousness. *PloS One*, *7*(5), e37238.

Owen, A. M., Coleman, M. R., Boly, M., Davis, M. H., Laureys, S., & Pickard, J. D. (2006). Detecting Awareness in the Vegetative State. *Science*, *313*(9), 1402.

Pape, T.B., Heinemann, A. W., Kelly, J. P., Hurder, A. G., & Lundgren, S. (2005). A measure of neurobehavioral functioning after coma. Part I: Theory, reliability, and validity of the Disorders of Consciousness Scale. *Journal of Rehabilitation Research and Development*, *42*(1), 1–17.

Pape, T. .-B., Senno, R. G., Guernon, A., & Kelly, J. P. (2005). A measure of neurobehavioral functioning after coma. Part II: Clinical and scientific implementation. *Journal of Rehabilitation Research and Development*, *42*(1), 19–28.

Parisot, S., Ktena, S. I., Ferrante, E., Lee, M., Guerrero, R., Glocker, B., & Rueckert, D. (2018). Disease prediction using graph convolutional networks: Application to

Autism Spectrum Disorder and Alzheimer' s disease. *Medical Image Analysis*, *48*, 117–130.

Park, H.-J., & Friston, K. J. (2013). Structural and Functional Brain Networks: From Connections to Cognition. *Science*, *342(6158)*, 1–8.

Pedregosa, F., Weiss, R., & Brucher, M. (2011). Scikit-learn : Machine Learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830.

Penttonen, M., & Buzsáki, G. (2003). Natural logarithmic relationship between brain oscillators Natural logarithmic relationship between brain oscillators. *Thalamus and Related Systems*, *6*(4), 145–152.

Phelps, E. A. (2006). Emotion and Cognition: Insights from Studies of the Human Amygdala, (Miller 2003), 27–53.

Phillips, C., Bruno, M.-A., Maquet, P., Boly, M., Noirhomme, Q., Schnakers, C., … Laureys, S. (2011). "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients. *NeuroImage*, *56*(2), 797–808.

Posner, J. B., Saper, C. B., Schiff, N. D., & Plum, F. (2007). *Diagnosis of Stupor and Coma*.

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *NeuroImage*, *59*(3), 2142–2154.

Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Anne, K., Church, J. A., … Petersen, S. E. (2011). Functional network organization of the human brain. *Neuron*, *72*(4), 665–678.

Power, J. D., Plitt, M., Gotts, S. J., Kundu, P., Voon, V., Bandettini, P. A., & Martin, A. (2018). Ridding fMRI data of motion-related influences : Removal of signals with distinct spatial and physical bases in multiecho data. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(9), E2105–E2114.

Preller, K. H., Burt, J. B., Ji, J. L., Schleifer, C. H., Adkinson, B. D., Stampfli, P., … Anticevic, A. (2018). Changes in global and thalamic brain connectivity in LSD-induced altered states of consciousness are attributable to the 5- HT2A receptor. *eLife*, *7*, 1–31.

Purdon, P. L., Pierce, E. T., Mukamel, E. A., Prerau, M. J., Walsh, J. L., Wong, K. F. ., … Brown, E. N. (2013). Electroencephalogram signatures of loss and recovery of consciousness from propofol. *Proceedings of the National Academy of Sciences of*

the United States of America, *110*(12), E1142-51.

Raichle, M. E. (1983). Positron Emission Tomography. *Annual Review of Neuroscience*, *6*, 249–267.

Raichle, M. E. (2010). Two views of brain function. *Trends in Cognitive Sciences*, *14*(4), 180–90.

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*(2), 676–682.

Rees, G., Kreiman, G., & Koch, C. (2002). Neural Correlates of Consciousness in Humans. *Nature Reviews Neuroscience*, *3*, 1–10.

Rowley, G., & Fielding, K. (1991). Reliability and accuracy of the Glasgow Coma Scale with experienced and inexperienced users. *The Lancet*, *337*, 535–538.

Rubia, K., Smith, A. B., Brammer, M. J., & Taylor, E. (2003). Right inferior prefrontal cortex mediates response inhibition while mesial prefrontal cortex is responsible for error detection. *NeuroImage*, *20*(1), 351–358.

Rubia, K., Smith, A. B., Taylor, E., & Brammer, M. (2007). Linear age-correlated functional development of right inferior fronto-striato-cerebellar networks during response inhibition and anterior cingulate during error-related processes. *Human Brain Mapping*, *28*(11), 1163–77.

Rudolph, U., & Antkowiak, B. (2004). Molecular and neuronal substrates for general anaesthetics. *Nature Reviews Neuroscience*, *5*(9), 709–720.

Rudy, B., Fishell, G., Lee, S., & Hjerling-leffler, J. (2013). Three Groups of Interneurons Account for Nearly 100% of Neocortical GABAergic Neurons, *71*(1), 45–61.

Ryali, S., Supekar, K., Abrams, D. A., & Menon, V. (2010). Sparse logistic regression for whole brain classification of fMRI data. *NeuroImage*, *51*(2), 752–764.

Sacco, S., Altobelli, E., Pistarini, C., Cerone, D., Cazzulani, B., & Carolei, A. (2011). Validation of the Italian version of the Coma Recovery Scale-Revised (CRS-R). *Brain Injury*, 1–8.

Salimi-Khorshidi, G., Douaud, G., Beckmann, C. F., Glasser, M. F., Griffanti, L., & Smith, S. M. (2014). NeuroImage Automatic denoising of functional MRI data : Combining independent component analysis and hierarchical fusion of classi fi ers. *NeuroImage*, *90*, 449–468.

Scannell, J. W., Burns, G. A. P. C., Hilgetag, C. C., O'Neil, M. A., & Young, M. P. (1999). The connectional organization of the cortico-thalamic system of the cat The Connectional Organization of the Cortico-thalamic System of the Cat. *Cerebral Cortex*, *9*, 1047–3211.

Schiff, N. D., Rodriguez-Moreno, D., Kamal, A., Kim, K. H. S., Giacino, J. T., Plum, F., & Hirsch, J. (2005). fMRI reveals large-scale network activation in minimally conscious patients. *Neurology*, *64*(3), 514-523

Schnakers, C., Majerus, S., Giacino, J. T., Vanhaudenhuyse, A., Bruno, M.-A., Boly, M., … Laureys, S. (2008). A French validation study of the Coma Recovery Scale-Revised (CRS-R). *Brain Injury*, *22*, 786–792. https://doi.org/10.1080/02699050802403557

Schroter, M. S., Spoormaker, V. I., Schorer, A., Wohlschla, A., Czisch, M., Kochs, E. F., … Jordan, D. (2012). Spatiotemporal Reconfiguration of Large-Scale Brain Functional Networks during Propofol-Induced Loss of Consciousness. *Journal of Neuroscience*, *32*(37), 12832–12840.

Schrouff, J., Perlbarg, V., Boly, M., Marrelec, G., Boveroux, P., Vanhaudenhuyse, A., … Benali, H. (2011). Brain functional integration decreases during propofol-induced loss of consciousness. *NeuroImage*, *57*(1), 198–205.

Schultz, D. H., & Cole, M. W. (2016). Higher Intelligence Is Associated with Less Task-Related Brain Network Reconfiguration. *Journal of Neuroscience*, *36*(33), 8551–8561.

Seamans, J. (2008). Losing inhibition with ketamine. *Nature Chemical Biology*, *4*(2), 91–93.

Seel, R. T., Sherer, M., Whyte, J., Katz, D. I., Giacino, J. T., Rosenbaum, A. M., … Zalser, N. (2010). Assessment scales for disorders of consciousness: evidence-based recommendations for clinical practice and research. *Archives of Physical Medicine and Rehabilitation*, *91*(12), 1795–1813.

Setsompop, K., Kimmlingen, R., Eberlein, E., Witzel, T., Cohen-adad, J., Mcnab, J. A., … Wald, L. L. (2013). NeuroImage Pushing the limits of in vivo diffusion MRI for the Human Connectome Project. *NeuroImage*, *80*, 220–233.

Shewmon, D. A., Holmes, G. L., & Byrne, P. A. (1999). Consciousness in congenitally decorticate children : developmental vegetative state as self-fulfilling prophecy. *Developmental Medicine and Child Neurology*, *41*, 364–374.

Shiel, A., Horn, S. A., Wilson, B. A., Watson, M. J., Campbell, M. J., & McLellan, D. L.

(2000). The Wessex Head Injury Matrix (WHIM) main scale: a preliminary report on a scale to assess and monitor patient recovery after severe head injury. *Clinical Rehabilitation*, *14*, 408–416.

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*(5), 648–663.

Simonyan, K., & Zisserman, A. (2015). Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations* (1–14).

Skudlarski, P., Constable, R. T., & Gore, J. (1999). ROC Analysis of Statistical Methods Used in Functional MRI : Individual Subjects. *NeuroImage*, *9*, 311–329.

Smith, M. M., Citerio, G. G., & Kofke, W. A. (2016). *Oxford Textbook of Neurocritical Care*.

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., … Beckmann, C. F. (2009). 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*, 13040–13045.

Snoek, J., Rippel, O., Swersky, K., Kiros, R., Satish, N., Sundaram, N., … Adams, R. P. (2015). Scalable Bayesian Optimization Using Deep Neural Networks. In *Proceedings of the 32nd International Conference on Machine Learning*.

Sohal, V. S., Zhang, F., Yizhar, O., & Deisseroth, K. (2009). Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature*, *459*(7247), 698–702.

Sporns, O., Tononi, G., & Kotter, R. (2005). The Human Connectome : A Structural Description of the Human Brain. *PLoS Computational Biology*, *1*(4), e42.

Spreng, R. N. (2012). The fallacy of a "task-negative" network. *Frontiers in Psychology*, *3*(5), 1–5.

Spreng, R. N., Sepulcre, J., Turner, G. R., Stevens, W. D., & Schacter, D. L. (2013). Intrinsic architecture underlying the relations among the default, dorsal attention, and frontoparietal control networks of the human brain. *Journal of Cognitive Neuroscience*, *25*(1), 74–86.

Spreng, R. N., Stevens, W. D., Chamberlain, J. P., Gilmore, A. W., & Schacter, D. L. (2010). Default network activity, coupled with the frontoparietal control network, supports goal-directed cognition. *NeuroImage*, *53*(1), 303–317.

Spreng, R. N., Stevens, W. D., Viviano, J. D., & Schacter, D. L. (2016). Neurobiology of Aging Attenuated anticorrelation between the default and dorsal attention networks with aging : evidence from task and rest. *Neurobiology of Aging*, *45*, 149–160. https://doi.org/10.1016/j.neurobiolaging.2016.05.020

Stamatakis, E. A., Adapa, R. M., Absalom, A. R., & Menon, D. K. (2010). Changes in resting neural connectivity during propofol sedation. *PloS One*, *5*(12), e14224. https://doi.org/10.1371/journal.pone.0014224

Stamatakis, E. A., Orfanidou, E., & Papanicolaou, A. C. (2017). Functional Magnetic Resonance Imaging Physiological Processes Underlying the fMRI. In *The Oxford Handbook of Functional Brain Imaging in Neuropsychology and Cognitive Neurosciences* (pp. 1–25).

Stender, J., Gosseries, O., Bruno, M.-A., Charland-verville, V., Vanhaudenhuyse, A., Demertzi, A., … Laureys, S. (2014). Diagnostic precision of PET imaging and functional MRI in disorders of consciousness : a clinical validation study. *The Lancet*, *6736*(14), 8–16.

Stepan, C., Haidinger, G., & Binder, H. (2004). Prevalence of persistent vegetative state/apallic syndrome in Vienna. *European Journal of Neurology*, *11*, 461–466.

Supp, G. G., Siegel, M., Hipp, J. F., & Engel, A. K. (2011). Report Cortical Hypersynchrony Predicts Breakdown of Sensory Processing during Loss of Consciousness. *Current Biology*, *21*(23), 1988–1993.

Teasdale, G., & Jennet, B. (1974). Assessment of Coma and Impaired Consciousness. *The Lancet*, 81–84.

Teasdale, G., Maas, A., Lecky, F., Manley, G., Stocchetti, N., & Murray, G. (2014). The Glasgow Coma Scale at 40 years: standing the test of time. *The Lancet Neurology*, *13*, 844–854.

Thomas Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., … Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*, 1125–1165.

Tononi, G., Boly, M., Massimini, M., & Koch, C. (2016). Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience*, *17*(7), 450–461.

Traub, R. D., Jefferys, J. G. R., & Whittington, M. A. (1997). Simulation of Gamma Rhythms in Networks of Interneurons and Pyramidal Cells. *Journal of*

Computational Neuroscience, *150*, 141–150.

Tremblay, R., Lee, S., & Rudy, B. (2016). GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits. *Neuron*, *91*(2), 260–292.

Tresch, D. D., Sims, F. H., Duthie, E. H., Goldstein, M. D., & Lane, P. S. (1991). Clinical Characteristics of Patients in the Persistent Vegetative State. *Archives of Internal Medicine*, *151*, 930–932.

Tronel, C., Largeau, B., Joao, M., Ribeiro, S., Guilloteau, D., Dupont, A.-C., & Arlicot, N. (2017). Molecular Targets for PET Imaging of Activated Microglia : The Current Situation and Future Expectations. *International Journal of Molecular Sciences*, *18*, 1–22.

Tuch, D. S. (2004). Q-Ball Imaging. *Magnetic Resonance Imaging in Medicine*, *52*, 1358–1372.

Tung, A., Tadimeti, L., Atkins, P. M., Mion, L. C., Palmer, R. M., Slomka, J., & Mendelson, W. (2001). The Relationship of Sedation to Deliberate Self-Extubation. *Journal of Clinical Anesthesia*, *13*, 24–29.

Uddin, L. Q., Kelly, A. M. C., Biswal, B. B., Castellanos, F. X., & Milham, M. P. (2009). Functional Connectivity of Default Mode Network Components: Correlation , Anticorrelation , and Causality. *Human Brain Mapping*, *637*(1), 625–637.

Van Calster, L., Argembeau, A. D., Salmon, E., Peters, F., & Majerus, S. (2017). Fluctuations of Attentional Networks and Default Mode Network during the Resting State Reflect Variations in Cognitive States: Evidence from a Novel Resting-state Experience Sampling Method. *Journal of Cognitive Neuroscience*, *29*(1), 95–113.

van den Heuvel, M., & Sporns, O. (2011). Rich-Club Organization of the Human Connectome. *Journal of Neuroscience*, *31*(44), 15775–15786.

Van Dijk, K. R. A., Hedden, T., Venkataraman, A., Evans, K. C., Lazar, S. W., & Buckner, R. L. (2018). Intrinsic Functional Connectivity As a Tool For Human Connectomics : Theory , Properties , and Optimization. *Journal of Neurophysiology*, *2138*, 297–321.

van Erp, W. S., Lavrijsen, J. C. M., van de Laar, F. A., Vos, P. E., & Laureys, S. (2014). The vegetative state / unresponsive wakefulness syndrome : a systematic review of prevalence studies. *European Journal of Neurology*, *21*, 1361–1368.

van Gaal, S., & Lamme, V. A. F. (2012). Unconscious High-Level Information Processing : Implication for Neurobiological Theories of Consciousness. *The*

Neuroscientist, *18*(3), 287–301.

Van Keulen, S. G., & Burton, J. H. (2003). Myoclonus Associated With Etomidate for ED Procedural Sedation and Analgesia. *American Journal of Emergency Medicine*, *21*(7), 556–558.

Vande Voort, J. L., Swintak, C. C., Wall, C. A., & Rasmussen, K. G. (2013). Methohexital-Induced Seizures During Electroconvulsive Therapy. *Journal of Electroconvulsive Therapy*, *29*(1), 4–5.

Vanhaudenhuyse, A., Noirhomme, Q., Tshibanda, L., Bruno, M.-A., Boveroux, P., Schnakers, C., … Boly, M. (2010). Default network connectivity reflects the level of consciousness in non-communicative brain-damaged patients. *Brain*, *133*, 161–171.

Varshney, L. R., Chen, B. L., Paniagua, E., Hall, D. H., & Chklovskii, D. B. (2011). Structural Properties of the Caenorhabditis elegans Neuronal Network. *PLoS Computational Biology*, *7*(2), e1001066.

Váša, F., Bullmore, E. T., & Patel, A. X. (2018). Probabilistic thresholding of functional connectomes: Application to schizophrenia. *NeuroImage*, *172*(8), 326–340.

Vatansever, D., Manktelow, A. E., Sahakian, B. J., Menon, D. K., & Stamatakis, E. A. (2016). Cognitive Flexibility: A Default Network and Basal Ganglia Connectivity Perspective. *Brain Connectivity*, *6*(3), 201–207.

Vatansever, D., Menon, D. K., Manktelow, A. E., Sahakian, B. J., & Stamatakis, E. A. (2015a). Default Mode Dynamics for Global Functional Integration. *Journal of Neuroscience*, *35*(46), 15254–62.

Vatansever, D., Menon, D. K., Manktelow, A. E., Sahakian, B. J., & Stamatakis, E. A. (2015b). Default mode network connectivity during task execution. *NeuroImage*, *122*, 96–104.

Vatansever, D., Menon, D. K., & Stamatakis, E. A. (2017). Default mode contributions to automated information processing. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(48), 12821–12826.

Vincent, J. L., Patel, G. H., Fox, M. D., Snyder, A. Z., Baker, J. T., Essen, D. C. Van, … Raichle, M. E. (2007). Intrinsic functional architecture in the anaesthetized monkey brain, *447*(5), 83-86.

Volicer, L., Berman, S. A., Cipolloni, P. B., & Mandell, A. (1997). Persistent Vegetative State in Alzheimer Disease: Does It Exist? *Archives of Neurology*, *54*, 1382–1384.

Walshe, T. M., & Leonard, C. (1985). Persistent Vegetative State Syndrome to Include Chronic Disorders. *Archives of Neurology*, *42*, 1045–1047.

Wang, B., Pourshafeie, A., Zitnik, M., Zhu, J., Bustamante, C. D., & Leskovec, J. (2018). Network enhancement as a general method to denoise weighted biological networks. *Nature Communications*, *9*, 1–8.

Weissenbacher, A., Kasess, C., Gerstl, F., Lanzenberger, R., Moser, E., & Windischberger, C. (2009). Correlations and anticorrelations in resting-state functional connectivity MRI: A quantitative comparison of preprocessing strategies. *NeuroImage*, *47*(4), 1408–1416.

Whitaker, K. J., Vertes, P., Romero-Garcia, R., Vasa, F., Moutoussis, M., Prabhu, G., … Bullmore, E. T. (2016). Adolescence is associated with genomically patterned consolidation of the hubs of the human brain connectome. *Proceedings of the National Academy of Sciences of the United States of America*, *113*(32), 9105–9110.

White, J. G., & Southgate, E. (1986). The structure of the nervous system of the nematode caenorhabditis elegans, *340*, 1–340.

Whitfield-Gabrieli, S., & Nieto-Castañón, A. (2012). Conn: A Functional Connectivity Toolbox for Correlated and Anticorrelated Brain Networks. *Brain Connectivity*, *2*(3), 125–141.

Whittington, M. A., & Traub, R. D. (2003). Interneuron Diversity series : Inhibitory interneurons and network oscillations in vitro. *Trends in Neurosciences*, *26*(12), 676–682.

Whittington, M. A., Traub, R. D., & Jefferys, J. G. R. (1995). Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation. *Nature*, *373*(16), 612–615.

Wu, C. W., Gu, H., Lu, H., Stein, E. A., Chen, J.-H., & Yang, Y. (2008). Frequency specificity of functional connectivity in brain networks. *NeuroImage*, *42*, 1047–1055.

Yeh, F.-C., Verstynen, T. D., Wang, Y., Fernández-Miranda, J. C., & Isaac, W.-Y. (2013). Deterministic Diffusion Fiber Tracking Improved by Quantitative Anisotropy. *PloS One*, *8*(11), 1–16.

Yeh, F.-C., Wedeen, V. J., & Tseng, W.-Y. I. (2010). Generalized q -Sampling Imaging. *IEEE Tansactions on Medical Imaging*, *29*(9), 1626–1635. https://doi.org/10.1109/TMI.2010.2045126

Yip, G. M. S., Chen, Z.-W., Edge, C. J., Smith, E. H., Dickinson, R., Hohenester, E., … Franks, N. . (2013). A propofol binding site on mammalian GABAA receptors identified by photolabeling. *Nature Chemical Biology*, *9*(11), 715–20.

Zalesky, A., Fornito, A., & Bullmore, E. T. (2010). Network-based statistic: Identifying differences in brain networks. *NeuroImage*, *53*(4), 1197–1207.

Zhang, M., Cui, Z., Neumann, M., & Chen, Y. (2018). An End-to-End Deep Learning Architecture for Graph Classification. In *The Thirty-Second AAAI Conference on Artificial Intelligence* (pp. 4438–4445).