

Diffusion MRI Analysis of TMS effects on Cognitive Pain Processing and Descending Pain Modulation in Chronic Knee Osteoarthritis

Max C. Wang

Supervisors: Dr. Dorothee Auer

Dr. Stamatios Sotiropoulos

Dr. Rob Dineen

Thesis submitted to the University of Nottingham
for the degree Master of Research

July 2024

Table of Contents

Abstract-----	i
Acknowledgements -----	ii
Abbreviations -----	iii
I. Introduction and Background -----	1
1. Chronic pain -----	1
1.1 The burden of chronic pain -----	1
1.2 Peripheral nociception in acute pain-----	1
1.4. Descending pain modulation-----	4
1.5. Chronic pain and abnormal pain processing-----	6
1.6 Chronic pain in knee osteoarthritis -----	7
2. MRI imaging of structural neuroplasticity-----	7
2.1. Basic MRI principles -----	8
2.2. dMRI contrast and acquisition-----	9
2.3. Diffusion tensor imaging (DTI)-----	14
2.4. Diffusion kurtosis imaging (DKI) -----	15
2.5. Neurite orientation dispersion and density imaging (NODDI) -----	16
3. Modulating the cortex with TMS -----	18
3.1. Therapeutic neuromodulation methods in chronic pain: why TMS?-----	18
3.2. Basic TMS principles-----	20
3.3. Connectivity-guided intermittent theta-burst stimulation-----	22
3.4. TMS in chronic pain -----	22
4. BoostCPM: A clinical trial using iTBS in patients with chronic knee OA -----	23
II. Methodology-----	25
1. Study design -----	25
1.1. Participants -----	25
1.2. iTBS protocol -----	25
1.3. MRI protocol-----	26
2. MRI processing and analysis-----	27
2.1. Software and setup -----	27
2.2. Structural processing and QC -----	27
2.3. dMRI processing and QC -----	31
2.4. ROI definitions -----	35
2.5. Extraction of dMRI outcomes -----	38
3. Statistical testing-----	38
III. Results-----	40
1. Quality control of MRI data and analysis -----	40
1.1. Data exclusion and harmonization -----	40
1.2. QC of structural preprocessing -----	40

1.3. QC of dMRI preprocessing-----	42
1.4. Comparison of Freesurfer and Glasser ROIs-----	45
2. Structural neuroplasticity-----	46
2.1. Neuroplastic changes: MD-----	46
2.2. Neuroplastic changes: MK-----	50
2.3. Neuroplastic changes: NODDI -----	54
IV. Discussion and Conclusion-----	64
1. Robustness of results -----	64
1.1. Pipeline variability and QC considerations -----	64
1.2. Validity of quantitative dMRI results-----	65
2. Interpretation and implications -----	65
2.1. Interpretation of quantitative results-----	65
2.2. Clinical implications-----	67
3. Choice of methods and future directions -----	68
3.1. Potential alternative methods-----	68
3.2. Future analyses-----	69
4. Conclusion -----	69
V. Appendices -----	71
Appendix A: Understanding the off-diagonal elements in the diffusion tensor-----	71
Appendix B: TMS protocol parameters-----	72
VI. Bibliography -----	73

Abstract

It is increasingly recognized that chronic pain is characterized by abnormalities in the neural circuits that process the pain signal. The anterior insula (AI) and anterior cingulate cortex (ACC) are two key hubs of high-level attentional processing that have been observed to be altered in chronic pain. Furthermore, the descending pain modulation system (DPMS) that exerts top-down control over afferent nociceptive signals can exhibit maladaptive dynamics in chronic pain. Transcranial magnetic stimulation (TMS) has been proposed as a non-invasive therapy for inducing normalization of aberrant brain circuits. A randomized sham-controlled clinical trial (BoostCPM) was conducted to evaluate the efficacy of an accelerated TMS protocol for treating chronic pain in a cohort of patients with knee osteoarthritis. Diffusion MRI (dMRI) was collected at baseline and after TMS treatment to assess for neuroplastic changes in key pain processing centers.

This thesis presents an analysis of the dMRI data from the BoostCPM trial. The details of this image analysis were defined a priori in a documented plan uploaded to the University of Nottingham's Research Data Repository (<http://doi.org/10.17639/nott.7388>). The dMRI data was processed through an extensive pipeline and closely inspected for quality control at every step. Minor errors in data processing were noted and discussed alongside final results. Grey matter hubs of cognitive pain processing and the DPMS were chosen as regions of interest. Three advanced dMRI models were applied, generating five quantitative indices of neural microstructure. Statistical tests were performed for comparisons within-group (pre- vs. post-TMS) and between-group (active vs. sham intervention). Six significant results were obtained out of 135 tests; none of these survived p-value correction for multiple tests. Interpretation the results within the context of the analysis methods used and plausible neurobiological mechanisms tends to suggest that the six significant test results are spurious. It is most reasonably concluded that accelerated TMS does not induce structural neuroplasticity in grey matter pain centers within this cohort of knee osteoarthritis patients.

Acknowledgements

This space is for expressing gratitude for the people and organizations who have helped me with this work—research is never an individual endeavor. I would like to thank my primary supervisor Dr. Dorothee Auer for her mentorship and guidance, for encouraging my independent lines of thinking, and for providing structure when necessary. I would also like to thank my supervisors Dr. Stamatios Sotiropoulos and Dr. Rob Dineen for offering their insight and resources throughout the project. I am indebted to Dr. Duncan Hodgkinson who was a quotidian presence throughout the year and a close mentor of my work. Thank you to Dr. Christopher Tench who provided guidance on the statistical aspects of this project.

I am grateful for the students and postdocs at the Sir Peter Mansfield Imaging Center (SPMIC) at Queen's Medical Centre who entertained my frequent and often wandering questions with patience. In particular, thank you to Dr. Alireza Mohammadi-Nejad and Dr. Martin Craig for their guidance on the technical aspects of this project, and to Jenna Hanmer and Francesco D'Antonio for their friendship, camaraderie, and support.

I appreciate the support from the University of Nottingham and its resources, in particular the high performance computing clusters (Augusta and Ada) that were integral to this project.

Thank you to the US-UK Fulbright Commission for granting me the privilege and means to pursue this work in Nottingham.

Abbreviations

- **A/P:** Anterior-posterior
- **ACC:** Anterior cingulate cortex
 - (r/l)AntACC: (Right/Left) anterior cingulate cortex
 - (r/l)MidAntACC: (Right/Left) mid-anterior cingulate cortex
- **ADC:** Apparent diffusion coefficient
- **AI:** Anterior insula
 - rAI: Right anterior insula
- **AKC:** Average kurtosis coefficient
- **ANTs:** Advanced Normalization Tools
- **cg-iTBS:** Connectivity-guided iTBS
- **CPM:** Conditioned pain modulation
- **CSF:** Cerebrospinal fluid
- **DBS:** Deep brain stimulation
- **DKI:** Diffusion kurtosis imaging
- **DLPFC:** Dorsolateral prefrontal cortex
 - lDLPFC: Left DLPFC
- **dMRI:** Diffusion MRI
- **DNIC:** Diffuse noxious inhibitory control
- **DPMS:** Descending pain modulation system
- **DTI:** Diffusion tensor imaging
- **EMG:** Electromyography
- **EPI:** Echo planar imaging
- **FA:** Fractional anisotropy
- **FDR:** False discovery rate
- **FLAIR:** Fluid-attenuated inversion recovery
- **fMRI:** Functional MRI
- **FOV:** Field of view
- **FSL:** FMRIB Software Library
- **FUS:** Focused ultrasound stimulation
- **GM:** Grey matter
- **ICVF:** Intracellular volume fraction
- **ISOVF:** Isometric volume fraction
- **iTBS:** Intermittent theta-burst stimulation
- **M1:** Primary motor cortex
- **MD:** Mean diffusivity
- **MEP:** Motor-evoked potential
- **MK:** Mean kurtosis
- **MNI:** Montreal Neurological Institute atlas
- **MRI:** Magnetic resonance imaging
- **NIBS:** Non-invasive brain stimulation
- **NODDI:** Neurite orientation dispersion and density imaging
- **OA:** Osteoarthritis
- **ODI:** Orientation dispersion index
- **P/A:** Posterior-anterior
- **PAG:** Periaqueductal grey
- **PFC:** Prefrontal cortex
 - mPFC: medial prefrontal cortex
 - dmPFC: dorsomedial prefrontal cortex
 - (r/l)dmPFC: (Right/Left) dmPFC
 - vmPFC: ventromedial prefrontal cortex
 - (r/l)vmPFC: (Right/Left) vmPFC
- **PGSE:** Pulsed gradient spin echo
- **QC:** Quality control
- **QUAD:** Quality assessment for dMRI (FSL software)
- **RF:** Radiofrequency
- **rMT:** Resting motor threshold
- **ROI:** Region of interest
- **rs-fMRI:** Resting state fMRI
- **rTMS:** repetitive TMS
- **RVM:** Rostral ventromedial medulla
- **S1:** Primary somatosensory cortex
- **S2:** Secondary somatosensory cortex
- **SN:** Salience network
- **SNR:** Signal-to-noise ratio
- **tACS:** Transcranial alternating current stimulation
- **tDCS:** Transcranial direct current stimulation
- **TE:** Time to echo
- **TMS:** Transcranial magnetic stimulation
- **TR:** Time to repetition
- **WM:** White matter

I. Introduction and Background

1. Chronic pain

1.1 The burden of chronic pain

Chronic pain, a condition defined as continued or recurrent pain for more than 3 months [Treede 2019], affects an estimated 10-25% of people globally [Goldberg 2011]. The societal burden of chronic pain is consistently demonstrated to be a leading issue across countries. For example, the disorder has been estimated to cause the greatest loss in quality-adjusted life years in Singapore [Abdin 2020] and account for greater economic losses than heart disease and cancer in the United States [Gaskin 2012]. In the United Kingdom, the prevalence of chronic pain among adults has been estimated at 35-50% [Fayaz 2016].

Despite the prevalence and burden of chronic pain, current treatments are insufficient to manage the breadth of conditions. Opioids are commonly prescribed to manage the condition, but they demonstrate limited analgesic efficacy and are highly addictive. Opioids also carry associated risks for overdoses and other serious medical complications [Chou 2015]. Given the global burden of chronic pain and the significant shortcomings of current clinical treatment regimens, new analgesic therapies for managing the condition are urgently necessary.

1.2 Peripheral nociception in acute pain

To understand chronic pain, it is necessary to first establish the neuronal mechanics of acute pain. The acute pain pathway begins in the peripheral nervous system with the firing of primary nociceptors—sensory neurons that respond differentially to stimuli capable of causing tissue injury [Julius 2001]. The noxious stimulus can be mechanical, thermal, or chemical in nature, and indeed many nociceptors are polymodal (Figure 1) [Julius 2001]. The axons of afferent nociceptors are either thinly myelinated A δ fibers (responsible for fast, sharp pain) or unmyelinated C fibers (responsible for longer, diffuse pain).

The prominent ascending nociceptive pathway is the spinothalamic tract (Figure 1). This pathway begins with first-order peripheral nociceptors entering the central nervous system through the dorsal horn of the spinal cord, where they synapse onto second-order neurons. The dorsal horn is organized in laminae, with A δ and C fibers synapsing in laminae I and II [Todd 2010]. The dorsal horn laminae also contain a complex circuitry of interneurons which integrate and modulate the nociceptive signal. Among these are wide dynamic range neurons which respond to nociceptive and non-nociceptive neurons and generate an integrated response [Price 2003]. After processing in the dorsal horn, second-order neurons decussate to the anterolateral quadrant of the contralateral spinal cord, then ascend to the central lateral and ventral posterior lateral nuclei in the thalamus [Willis 1997]. Downstream processing of this sensory information then involves the primary (S1) and secondary (S2) somatosensory cortices, as well as the insula and anterior cingulate cortex [Fields 2004].

This pathway from primary afferent nociceptor to spinothalamic tract is a useful but simplified model of peripheral nociception. It should be noted that while the spinothalamic tract is a prominent pathway for peripheral nociception, there are other pain pathways that ascend in the ventrolateral (e.g. spinomesencephalic tract) and dorsal (e.g. spinocervicothalamic tract) quadrants of

the spinal cord [Willis 1997]. Additionally, although primary nociceptors can be broadly defined as neurons responding to noxious stimulus, recent RNA sequencing studies have found tens of genetically diverse sensory neurons capable of generating pain [Lindsay 2021]. Finally, the dorsal horn circuits integrating afferent nociceptive and non-nociceptive signals involve many heterogeneous populations in neurons with incompletely elucidated interactions [Todd 2010]. A full review of the nuances of peripheral nociception is beyond the scope of this thesis.




Figure 1. The acute pain pathway via spinothalamic tract. Primary nociceptors in the peripheral nervous system relay signals to second order neurons in the dorsal horn of the spinal cord. Second order neurons decussate and ascend in the anterolateral quadrant, then synapse onto third order neurons in the thalamus. Third order neurons relay the signal to the primary somatosensory cortex. Lower box: nociceptors can be polymodal, with ion channels sensitive to mechanical, chemical, and thermal stimuli. Upper box: axons of primary nociceptors are either thin, myelinated Aδ fibers or unmyelinated C fibers.

1.3. Cognitive and affective pain

Pain is not an exclusively sensory-discriminative experience. It is a multisensory phenomenon that heavily involves cognitive-affective components [Wiech 2016]. The insufficiency of sensory nociception alone to describe pain is apparent when considering a clinical condition such as fibromyalgia, where pain arises without clear evidence of ongoing tissue damage and central processing is considered to be abnormally pro-nociceptive. Cognitive and affective processing can also lead to potent analgesic effects as in the cases of placebo and attentional analgesia.

There are several proposed frameworks for conceptualizing the cognitive-affective dimensions of pain. A basic approach is to deconstruct the experience of pain into three sequential components: sensory pain (how much pain do you feel?), pain unpleasantness (how much does the pain bother you?), and secondary affect (how does the pain affect your sense of homeostasis?) [Price 2000]. Recent evidence from small mammal studies lends a strong biological basis to this conceptual breakdown—a subset of neurons in the basolateral amygdala of rats has been shown to encode pain unpleasantness and pain-avoidant behavior in a manner dissociable from sensory pain perception [Corder 2019]. Another approach is to view pain through the lens of salience processing. Salience processing refers to the central brain mechanisms for filtering and identifying important stimuli from the constant afferent stream of sensory information [Legrain 2011]. In the context of pain, it is proposed that increased attention to a noxious stimulus may amplify pain, and abnormal salience processing may contribute to chronic pain [Borsook 2013]. A “salience network” (SN) of brain regions has been identified [Seeley 2007], and two key nodes of the SN are consistently implicated in affective pain processing: the anterior insula (AI) and anterior cingulate cortex (ACC) (Figure 2). Experimental evidence shows that these areas are differentially active when modulating pain unpleasantness while holding peripheral pain constant [Price 2000]. Anatomical connectivity also implicates these regions in central pain processing, as ascending nociceptive pathways project to the AI and ACC both directly from the thalamus and after processing in S1 and S2 (Figure 2) [Price 2000]. Finally, cognitive processing of pain may be conceptualized through a computational model of “predictive coding” emphasizing the importance of expectation [Wiech 2016]. Predictive coding models offer specific structure to the framework salience processing whereby stimuli congruent with expectation are amplified, while those that do not fit expectation are diminished [Wiech 2016]. The phenomena of placebo analgesia and diminished effect of opioids when administered without subject knowledge readily fit predictive coding models.




Figure 2. Nociceptive processing involves the ACC and AI, two key nodes of the SN. Ascending signals from the spinothalamic tract are processed in the AI and ACC via direct connections from the thalamus and indirect connections via the S1 and S2. ACC: anterior cingulate cortex; AI: anterior insula; S1: primary somatosensory cortex; S2: secondary somatosensory cortex; SN: salience network.

In summary, cognitive and affective processing of nociceptive input can lead to potent pro- and anti-pain effects and is strongly mediated through the AI and ACC, two key nodes of the SN. It should be noted at this point that the brain regions involved in sensory-discriminative and cognitive-affective processing are not entirely distinct but are instead overlapping and highly interconnected [Wiech 2016].

1.4. Descending pain modulation

In addition to the afferent nociceptive pathway discussed in section 1.3. and cognitive-affective central processing described in section 1.4., there is a descending efferent system which can modulate the action of nociceptors in the dorsal horn of the spinal cord. This modulation is bidirectional, capable of both enhancing and inhibiting pain [Fields 2004]. Recent advances have elucidated the specific pathways and biochemistry of this descending control system. It is now established that the system consists of opioidergic and serotonergic projections from the periaqueductal grey (PAG) to the rostral ventromedial medulla (RVM) to the dorsal horn of the spinal cord [Fields 2004]. Higher control of the PAG is mediated by multiple cortical and subcortical regions including the amygdala, ACC, orbitofrontal cortex, and prefrontal cortex (PFC) [Fields 2004, Lindsay 2021]. Together, the inputs to PAG and its descending projections to RVM and the spinal cord can be conceptualized as a “descending pain modulation system” (DPMS). This section will review the anatomical and functional characteristics of the DPMS (Figure 3).




Figure 3. The descending pain modulation system (DPMS). The DPMS facilitates top-down control of ascending nociception through the PAG and RVM. Output from the RVM synapses in the dorsal horn of the spinal cord to modulate incoming nociceptive signals. The PAG is mediated by cortical structures like the ACC and mPFC, and subcortical structures like the amygdala. The top-down pathway is shown in blue, while the bottom-up pathway is shown in green. ACC: anterior cingulate cortex; mPFC: medial prefrontal cortex; PAG: periaqueductal grey; RVM: rostral ventromedial medulla.

Anatomically, the PAG is a small region of grey matter which sits adjacent to the mesencephalic aqueduct in the midbrain of the brainstem. Its shape is an open cylinder approximately 14mm in length and 4-5mm in diameter, and it can be differentiated into four columns with distinct afferent and efferent connectivity [Linnman 2012]. The dorsomedial (dm), lateral, and ventrolateral (vl) columns project to lower brainstem regions including the RVM, while the dorsolateral (dl) column projects to the midbrain and pons [Linnman 2012]. Early work in rats showed that electrical stimulation of the PAG leads to a decrease in pain-related behaviors [Reynolds 1969], and analogous direct stimulation in humans can induce analgesia in a subset of patients [Bittar 2005]. PAG stimulation produces analgesia by activating its efferent connections to RVM [Ossipov 2010]. The RVM is a collection of nuclei in the medullary portion of the brainstem and represents the final and most inferior relay point for descending pain modulation signals in the brain. This region contains two functionally distinct populations of neurons that highlight the dichotomy of descending modulation—these are termed “on-cells” and “off-cells.” The firing patterns of these two classes of neurons are diametrically opposed during painful behaviors: on-cells are activated, while off-cells are inhibited [Ossipov 2010]. More generally, firing of on-cells elicits pain-facilitating descending modulation, while firing of off-cells elicits pain-inhibiting modulation [Ossipov 2010]. In line with this functional specification, pharmacological activation of RVM with opioid injection enhances the firing of off-cells, inhibits on-cells, and leads to potent analgesia [Heinricher 1994]. The analgesic effect of PAG stimulation is notably blocked by administering naloxone (a μ -opioid receptor antagonist) to the RVM, demonstrating that descending pain inhibition is mediated through an endogenous opioidergic pathway [Fields 2004]. On- and off-cells synapse onto interneurons in the substantia gelatinosa in the dorsal horn of the spinal cord, where they modulate the activity of ascending nociception via serotonergic, GABAergic, and opioidergic mechanisms [Lubejko 2022].

The PAG receives input from an array of cortical and subcortical structures. Given the varied functions of PAG (e.g. vocalization, micturition, thermoregulation), only a subset of these connections is strongly linked to descending pain modulation [Benarroch 2012]. Cortical modulation of PAG is mainly mediated through the PFC [Ong 2019]. Retrograde tracing studies of the PAG localize cortical afference mainly in the medial PFC (mPFC), and this mPFC to PAG connection has been strongly linked to cortical control of the DPMS in animal models [Ong 2019, Huang 2019]. Recently, a specific set of connections from the basolateral amygdala to mPFC to vlPAG was shown to be altered in neuropathic pain in a way that inhibited descending anti-pain modulation via the RVM [Huang 2019]. The ACC has also been shown to project directly to PAG, and stimulation of ACC can also modulate descending pain control [Urien 2019]. Subcortically, the central nucleus of the amygdala has direct projections to the lateral and ventrolateral columns of PAG [Lindsay 2021]. Optogenetic stimulation of this connection has been shown to modulate pain-related behaviors in animal models, implicating this pathway in descending control [Lindsay 2021]. It is important to note that other cortical and subcortical regions are also capable of modulating the activity of PAG even in the absence of direct connections. Such indirect modulation can occur via the direct connections of mPFC, ACC, and CeA discussed above. For example, while the dorsolateral PFC (DLPFC) does not project directly to PAG, it is heavily connected with mPFC and can therefore influence descending pain modulation through this indirect connection [Ong 2019]. Thus, the complete picture of afferent modulation of PAG is highly complex, but the most prominent direct connections have been highlighted here.

1.5. Chronic pain and abnormal pain processing

Chronic pain is accompanied by abnormalities in pain processing on the cellular and functional levels. In other words, chronic pain is not simply caused by persistent activation of nociceptors, but rather is characterized by aberrations in how the pain signal is processed in the central nervous system. Since acute pain is multidimensional, so too are the mechanisms of chronic pain. This section will review evidence of abnormalities that contribute to chronic pain occurring in ascending nociception, cognitive-affective processing, and descending modulation.

Neurons in the ascending nociceptive pathway can be sensitized under certain conditions to exhibit enhanced excitability, thereby enhancing the pain signal [Basbaum 2009]. Sensitized neurons exhibit lower thresholds for activation and greater evoked potentials [Baron 2013]. Sensitization can occur peripherally in primary nociceptors or centrally in second-order spinal cord neurons. Peripheral sensitization can follow disease or injury, where chronic inflammation creates a biochemical environment that increases the excitability of primary nociceptors, contributing to a chronic pain state [Basbaum 2009]. Central sensitization in contrast, occurs in second-order neurons and interneurons of the spinal dorsal horn that process primary nociceptive signals, but not in the spinothalamic tract axons themselves [Woolf 2000]. This heightened state of central responsiveness can be transiently induced by strong activation of peripheral nociceptors [Woolf 2011]. If peripheral nociceptive input is sustained, central sensitization may persist and contribute to chronic pain [Baron 2013]. On the biomolecular level, central sensitization is mediated by strengthened glutamatergic connections between primary nociceptors and second-order spinal cord neurons, downregulation of GABAergic and glycinergic inhibitory interneurons in the dorsal horn, and activation of microglial cells [Basbaum 2009]. Clinically, central sensitization leads to hyperalgesia and allodynia. Hyperalgesia refers to an increased sensitivity to noxious stimuli, and allodynia refers to activation of the pain pathway by a typically innocuous stimuli (e.g. by a mechanoreceptor) [Woolf 2011]. Hyperalgesia can be a sign of peripheral or central sensitization, while allodynia suggests a neuropathic component to pain that includes but often extends beyond central sensitization.

Abnormalities in salience processing have also been proposed to be associated with pain chronification. Since the salience one assigns to a painful stimulus can dramatically alter the perceived intensity of pain, altered salience processing may contribute to the chronic pain state [Borsook 2013]. Animal models of chronic pain have demonstrated synaptic potentiation of the AI and ACC, two key brain regions that modulate salience processing [Tan 2021]. In humans, neuroimaging studies have shown structural and functional aberrations in these regions. Studies of the AI in chronic pain patients have indicated decreased grey matter volume, altered connectivity with nodes of other major resting-state brain networks, and increased activity when viewing pain-related words when compared to healthy controls [Borsook 2013, Lu 2016, Cottam 2018, Muthulingam 2020, Xu 2022]. In one study, altered functional connectivity of the AI was correlated with pain severity and accompanied by structural breakdown of white matter tracts between AI and ACC [Xu 2022]. These imaging studies suggest that cortical reorganization of salience processing regions may occur in chronic pain. Coarser evidence suggesting the importance of salience comes from post-cingulotomy patients who report being unbothered by their pain [Borsook 2013]. A final point to consider is that chronic pain is often comorbid with major depression and other affective disorders [Tunks 2008]. These comorbidities are also associated with altered activity in the insular and cingulate cortices, which may contribute to abnormal salience processing of pain [Borsook 2013].

Altered dynamics of the descending control system involving PAG and RVM may also contribute to chronic pain. It is proposed that healthy individuals exhibit a balance of inhibitory and facilitatory control in their DPMS, while chronic pain states may arise if descending inhibition fails [Ossipov 2014]. The inhibitory capability of the DPMS to suppress persistent nociception is evidenced by animal studies where hyperalgesia is elicited in injured rats only after suppression of the RVM [Ossipov 2014]. An experimental paradigm to measure descending inhibition in animals is diffuse noxious inhibitory control (DNIC), where a painful stimulus applied at a remote location inhibits neurons in the dorsal horn of the spinal cord [Le Bars 1979]. The analogous experimental procedure in humans is known as conditioned pain modulation (CPM), where subjective pain ratings to one stimulus are measured before and during application of a second painful stimulus at a remote location [Yarnitsky 2010]. It is important to note here that the DNIC effect is mediated through the caudal medulla in addition to the RVM, and the specific neurobiology of the CPM effect in relation to DNIC is a topic of active debate [Le Bars 1992, Sirucek 2023]. With this caveat in mind, it has been consistently shown that DNIC and CPM are diminished in chronic pain states, suggesting a decreased efficacy of descending inhibition via the DPMS [Lewis 2012, Chen 2019]. Plasticity of RVM neurons in animals exposed to persistent pain supports this notion [Chen 2019].

Abnormalities in afferent processing (peripheral and central sensitization), cognitive-affective processing (aberrant salience processing), and descending modulation (loss of descending inhibition and CPM) can all contribute to the chronic pain state. Since chronic pain can arise from many different conditions, the specific contributions from each abnormality can vary widely and is specific to the etiology for each patient.

1.6 Chronic pain in knee osteoarthritis

The body of this thesis will study a cohort of patients with chronic pain secondary to knee osteoarthritis (OA). OA is a degenerative joint disease where cartilage is lost due to a combination of mechanical, metabolic, or endocrine factors. The knee is the joint that is most frequently affected by OA, and musculoskeletal pain is a common symptom of the disease [Michael 2010]. While the pain caused by knee OA has a specific localizable source, chronic pain caused by the condition is more complex than sustained afferent nociception. Due to chronic inflammation and the persistent presence of a nocifensive stimulus, patients with chronic pain in knee OA sometimes exhibit peripheral and central sensitization [Fingleton 2015]. Thus, while chronic pain in knee OA was classically thought to be entirely nociceptive, many knee OA patients exhibit sensitization and some meet established criteria for neuropathic pain [Dimitroulas 2014]. Recent neuroimaging has found that the subset of patients with neuropathic pain exhibit altered functional activity of the ACC and RVM [Soni 2019]. The psychological impact of chronic pain in knee OA is variable. Patients who focus more on their pain and its associated negative affect (pain catastrophizing) are likely to be more psychologically and functionally disabled by their condition [Somers 2009].

2. MRI imaging of structural neuroplasticity

Although abnormalities in the CNS are increasingly recognized to contribute to the chronic pain state, the underlying mechanisms of these abnormalities are still unclear. One possible driver of CNS abnormalities is maladaptive neuroplasticity: changes in the structure and function of neurons and brain networks that deviate from the healthy norm. This neuroplasticity can be studied in humans using imaging. Over the past few decades, magnetic resonance imaging (MRI) has become a

dominant modality for imaging the structure and function of the human brain in vivo. Compared to other common techniques for imaging structure and function, (e.g. CT and PET), MRI has two distinct advantages. First, it does not require ionizing radiation and therefore minimizes risks to the subject. Both CT and PET generate image contrast by recording the absorption of ionizing radiation through tissue, a process that carries a risk of mutating somatic DNA which may lead to cancer. In contrast, MRI relies on differences in magnetic properties between tissue types to generate image contrast, a process that does not inherently expose the subject to additional risks. A second advantage of MRI is its multimodal nature, meaning it is possible to image both structure and function with a wide variety of contrasts within the same scanning session. Its multimodal capabilities arise from the large parameter space of acquisition and processing methods within the MRI framework. Within this parameter space, this thesis will focus on the utility of MRI for imaging structural neuroplasticity with a technique known as diffusion MRI (dMRI). Previous work using dMRI has demonstrated structural variability in the DPMS across individuals, and that this variability correlates with levels of endogenous pain modulation [Stein 2012]. If structural integrity in the DPMS is predictive of pain sensitivity, then modulating structural features may be therapeutic in chronic pain conditions. This thesis builds on this existing work by examining whether structural indices measured by dMRI can be experimentally altered by a novel brain stimulation technique. This section will first summarize basic MRI principles to ground the theory of dMRI acquisition, then describe how dMRI data is processed and modeled to index biological parameters of interest.

2.1. Basic MRI principles

This section reviews the basic principles of MRI imaging. Note that this is not meant to be a formal description of MRI physics, but rather a general overview to introduce concepts relevant to later discussions of dMRI. At the most fundamental level, the MRI signal arises from precessing ^1H protons in a strong magnetic field, and the contrast within an image comes from varying ^1H magnetization dynamics within a sample. When a subject is exposed to a strong static magnetic field (B_0), protons within the subject exhibit a slight preference for aligning with the field [Hanson 2008]. The aggregate effect of this directional preference is a net magnetic moment in line with the field. The manipulation and movement of this net magnetic moment induces current in receive coils via Faraday's law of induction. This is the basis of the MRI signal.

When imaging biological tissue, the overwhelming majority of protons contributing to the MRI signal are from water. Each individual proton precesses at a frequency (ω) proportional to the magnetic field strength (B_0) according to the Larmor Equation.

$$\omega = \gamma \cdot B_0$$

γ is the gyromagnetic ratio, a physical constant distinct to each atomic nucleus (for ^1H , $\gamma = 42.58$ MHz/T). The precession of protons in a magnetic field is known as Larmor precession.

For the simplest MRI acquisition, the first step to generating an image is to tip the net magnetization vector 90° into the transverse plane using a resonant radiofrequency (RF) excitation pulse. When precessing protons are subject to an RF pulse applied at their Larmor frequency, a resonance effect is observed which tips the axis of their spins towards the transverse plane. This resonance excitation underlies the etymology of *magnetic resonance* imaging. After the RF pulse is turned off, ^1H spins realign with the original net magnetic field. The movement of the net magnetic moment back towards the original configuration, or “relaxation” of spins, induces a signal in the receive coils of the MRI scanner. The resulting signal reflects the dynamics of relaxation, which vary between tissue

types. Variations in proton density, magnetic susceptibilities, and local chemical environment between tissues lead to different magnetic relaxation properties, giving rise to MRI contrast. The three main types of MR contrast are proton density, T1, and T2. Proton density simply reflects the concentration of resonant protons independent of relaxation properties. T1 captures the relaxation of net magnetization away from the transverse plane and back into the longitudinal axis, while T2 captures signal decay due to in-plane dephasing of spins from local magnetic effects.

MRI images are composed of voxels (3D volumetric pixels). The spatial localization of MRI signal to a specific voxel in the imaged subject is achieved through slice selection, phase and frequency encoding, and k-space reconstruction. First, a slice of the sample (typically axial) can be selectively excited by the RF pulse by leveraging the application of a gradient field along the longitudinal axis. This gradient modulates the Larmor frequency according to position along the longitudinal axis. Thus, the RF pulse frequency can be adjusted to selectively obtain signal from target slices in the sample (recall that RF frequency must match the Larmor frequency for resonant excitation to occur). Within a slice, two additional gradients are applied during acquisition to encode the x and y positions of the signal. One gradient, known as the phase-encoding gradient, is pulsed for a short period of time to induce a spatially varying phase across the slice. The other gradient, known as the frequency-encoding gradient, is applied in the orthogonal in-plane direction during signal readout to induce a spatially varying frequency across the slice. The raw data that is acquired for each slice exists in a frequency domain known as k-space, where the magnitude at each location specifies the contribution of a particular spatial frequency to the final image. Taking the Fourier transform of k-space generates the MRI image.

2.2. dMRI contrast and acquisition

Diffusion is the aggregate observable effect of random Brownian motion from individual particles. Brownian motion is the random movement of a particle suspended in a free medium due to thermally driven molecular collisions, and it was first observed of pollen suspended in water by Robert Brown in 1827 [Brown 1828]. For an individual particle, Brownian motion may be modeled as a random walk, where the direction of movement at each time point is given at random due to unpredictable molecular collisions. When observing the aggregate effect of random Brownian motion on a collection of particles, particles tend to disperse from areas of high concentration to areas of low concentration. It should be emphasized, though, that diffusion constantly occurs even when no concentration gradient exists. Within a solution of uniform concentration, the flux of diffusion across any surface is equal and opposite. Thus, diffusion is an inherent process for all particles within a free medium at any temperature above absolute zero, although it is most easily observed in the context of equalizing concentration gradients. Of relevance to dMRI is the concept of the diffusion coefficient (D). In its purest form, D is an inherent property of particles that describes how easily particles diffuse. It is related to the temperature (T), Boltzmann constant (k_B), viscosity of the solution (μ), and radius of the particle (r) as formulated by Einstein [Einstein 1905].

$$D = \frac{k_B T}{6\pi\mu r}$$

However, in the more macroscopic context of dMRI, diffusion coefficients are approximations of the aggregate diffusion occurring within a voxel. Voxels in which diffusion is hindered (e.g. by biological barriers) have lower D values than voxels with free diffusion.

Diffusion captures valuable neurobiological information because the dynamics of water diffusion within the brain vary widely between the three major tissue types: grey matter (GM), white matter (WM), and cerebrospinal fluid (CSF) (Figure 4). Since the magnetization of protons in water forms the basis of the MRI signal in biological contexts, MRI contrasts can be generated based on the distinct diffusion properties of different tissues. Importantly, the diffusion of water is hindered by the phospholipid bilayer of cell membranes. CSF is purely fluid; therefore, water molecules are free to diffuse in all directions without hindrance (high D). In contrast, the cell bodies of neurons in GM restrict the diffusion of water molecules, although not in any specific direction (low D). Most interesting are the diffusion characteristics of WM, which consists of axon tracts connecting GM regions. The cell membranes of WM tracts are tubular and therefore allow diffusion of water molecules in one specific direction while hindering diffusion in other directions. The diffusion of water in WM tracts is therefore highly anisotropic (high D in some directions, very low D in others). A visual graphic of the diffusion characteristics of GM, WM, and CSF is shown in Figure 4.




Figure 4: 2-dimensional diffusion profiles of the three main tissue types in the brain. GM and CSF have isotropic diffusion profiles, while white matter has characteristically anisotropic diffusion due to the tubular structure of axons. Diffusion is unhindered in CSF, while cellular membranes in GM and WM decrease the diffusivity of water molecules. CSF: cerebrospinal fluid, GM: grey matter, WM: white matter.

The simplest way to demonstrate how diffusion contrasts can be achieved with MRI is through a basic pulsed gradient spin echo (PGSE) sequence [Stejskal 1965]. The pulse diagram for this sequence is shown in Figure 5. Shortly after a 90° RF pulse tips the spins into the transverse plane, a gradient is applied to dephase the spins. That is, as protons precess at different rates due to the field gradient, they acquire a relative phase dependent on their position along the gradient. After some time ($TE/2$), a 180° RF pulse is applied to induce a spin echo which can be read at time TE . Immediately following the inversion pulse, a gradient with the same duration and strength as the first pulsed gradient is applied. Since the precessing protons were inverted with a 180° pulse, this gradient now has an equal and opposite effect on the phase of proton. In other words, the gradient-induced phases acquired after the 90° pulse are canceled out by the second gradient pulse. However, protons only completely rephase if they are stationary. Any proton which has diffused to a different location between the two pulses will gain a slightly different amount of phase during the second pulse and thus will not completely rephase. The result is an attenuation of the MRI signal that is proportional to the dephasing of the proton, which in turn depends on the distance it has diffused. In a PGSE sequence, low signal in a voxel suggests freer diffusion compared to a voxel with high signal. When imaging the brain, the signal will be lowest in CSF where free diffusion occurs. Using this contrast, a map of apparent diffusion coefficient (ADC) can be calculated where high ADC values (e.g. in CSF) indicate freer diffusion compared to lower ADC values (e.g. in GM) [Le Bihan 1986]. Note that all dMRI acquisitions are based on a T2-weighted spin-echo since dephasing due to diffusion primarily affects the dynamics of transverse relaxation reflected in T2 images.

There are two crucial parameters that define dMRI acquisitions: the diffusion direction and b-value. The diffusion direction is the direction in which the pulsed gradient is applied, and this determines the axis through which apparent diffusion is weighted. To illustrate this concept, consider a rigid tube which only allows water to diffuse in the x-direction. If the pulsed gradient is in the x-direction, the spins will diffuse along the direction of the gradient, leading to dephasing and generating diffusion contrast. However, if the pulsed gradient is in the y-direction, no protons diffuse along the gradient, and therefore no dephasing (and no diffusion contrast) occurs. While it is possible to generate an ADC map with a single diffusion direction, more accurate modeling of diffusion dynamics in modern dMRI implementations requires multiple diffusion directions.

dMRI acquisitions are also defined by their b-value, a parameter which defines the strength of the diffusion weighting. From the following equation (derived from a formal description of PGSE dynamics outside the scope of this thesis), the b-value of a sequence depends primarily on two factors: the diffusion time (δ) and the strength of the pulsed gradient (G).

$$b = \gamma^2 G^2 \Delta^2 \left(\delta - \frac{\Delta}{3} \right)$$

Increasing the diffusion time (the time between the two pulsed gradients, δ) allows more time for spins to diffuse and dephase, leading to a more attenuated signal. Increasing the strength of the pulsed gradient (G) causes greater dephasing given the same distance of diffusion, also leading to a more attenuated signal. Thus, acquisitions with higher b-values have greater diffusion-related signal dropout, which generates images that are highly sensitive to diffusion but have poor signal-to-noise ratio. Note that an image with $b=0$ has no diffusion weighting and is simply T2-weighted. Since the contrast of all diffusion-weighted images ($b>0$) also reflects T2-weighting alongside diffusion effects, all dMRI acquisitions will acquire b_0 volumes to normalize the T2 contrasts. Stejskal and




Figure 5. Theory of a simple PGSE pulse sequence acquisition. Top: the PGSE pulse sequence diagram. A diffusion gradient with strength G and duration Δ is applied shortly after the first RF pulse of 90° . After time $TE/2$, an inversion RF pulse is applied, and an identical diffusion gradient is pulsed. This leads to rephasing of spins at time TE . Bottom: the phase effect on a diffusing proton. The inversion pulse reverses the phase acquired during the first diffusion gradient, therefore the net phase acquired is proportional to the distance diffused. PGSE: pulsed gradient spin echo; RF: radiofrequency.

Tanner derived the following equation which relates signal attenuation to b-value and the diffusion coefficient (D) [Stejskal 1965].

$$\ln\left(\frac{S(b)}{S_0}\right) = -bD$$

$S(b)$ is the signal in a diffusion-weighted image of b-value equal to b , D is the diffusion coefficient, and S_0 is the signal in a T2-weighted image ($b=0$).

It can be observed from this relation that higher diffusivity D will lead to greater attenuation of the signal $S(b)$. Also note that deriving D from the relation requires normalizing the signal $S(b)$ to S_0 , (normalizing for the inherent T2 contrast of the sequence).




Figure 6: K-space acquisition for echo planar imaging (EPI) occurs in one frequency encoding step; taking the Fourier transform of k-space results in an image that is heavily distorted in the phase-encoding direction. Reversing the direction of phase-encoding in k-space causes spatial distortions to occur in the opposite direction.

Modern dMRI sequences use an echo-planar imaging (EPI) readout, where k-space is acquired in a single excitation. The usage of EPI significantly decreases the amount of time required to acquire a single diffusion-weighted image, therefore enabling many diffusion directions and b-values to be collected in a short amount of time. EPI also reduces motion artifacts which can skew diffusion estimates by minimizing the window for motion to affect the scan. In contrast with basic readouts, where a single frequency-encoding gradient is applied during the readout (thus filling one line of k-space), EPI readouts quickly alternate the frequency-encoding gradient to fill multiple lines of k-space in one echo (Figure 6).

However, EPI images are highly susceptible to geometric distortions. Because the entirety of k-space is acquired during an extended application of a single frequency-encoding gradient, deviations from theoretical precession frequencies due to inhomogeneity accumulate throughout the acquisition. Distortions affect the phase-encoding direction disproportionately because it is acquired

with a lower bandwidth (i.e. successive points in the phase-encoding direction are acquired at a lower rate in k-space compared to the frequency-encoding direction). The lower bandwidth of the phase-encoding direction means that phase errors (due to deviations from the expected precession frequency) have more time to accumulate along this axis. Thus, the encoded spatial information in the phase-encoding direction is subject to error caused by deviations from the expected magnetic field. In practice, this leads to an observed stretching or compression of the image along the phase-encoding direction in regions of high magnetic field inhomogeneity (Figure 6).

2.3. Diffusion tensor imaging (DTI)

While a single diffusion-weighted image only provides information about apparent diffusion in one direction, combining information from multiple diffusion-weighted scans can elucidate detailed information about the overall diffusion characteristics in each voxel. Starting with a set of assumptions about the diffusive properties of neural tissue, one can mathematically formulate dMRI models for deriving quantitative diffusion metrics that describe microstructure. The simplest and most commonly used dMRI model is diffusion tensor imaging (DTI). In each voxel, the DTI model fits a diffusion tensor: a 3x3 matrix which describes the 3-dimensional diffusivity profile within the voxel. Conceptually, the diffusion tensor (D) is a 3-dimensional diffusion coefficient.

$$D = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix}$$

The 3 diagonal elements of the matrix (D_{xx} , D_{yy} , D_{zz}) describe the apparent diffusivity in each of the applied gradient directions (x, y, and z), while 3 off-diagonal elements (D_{xy} , D_{xz} , D_{yz}) describe the correlation between pairs of diffusion directions. A hypothetical case illustrating the importance of the off-diagonal elements of the matrix is described in Appendix A. Note that D_{xy} and D_{yx} both quantify correlation between the same pair of directions and are therefore equivalent (the same is true for D_{xz} with D_{zx} , and D_{yz} with D_{zy}). Thus, the diffusion tensor contains 6 independent components and therefore requires at least 6 diffusion-weighted scans in unique directions to estimate. A formal derivation of how to estimate the diffusion tensor from dMRI (first described by Bassar and colleagues in 1994) is out of the scope of this thesis, but it should be stated that the calculation of D is based on a generalization of the 1D Stejskal-Tanner relation into 3-dimensional space. For an applied gradient (G) in a direction with x, y, and z components, the equation relates the T2-normalized signal attenuation ($S(b)/S_0$) to a linear combination of b-values which depend on the gradient direction [Bassar 1994].

$$\ln\left(\frac{S(b)}{S_0}\right) = -\left(b_{xx}D_{xx} + b_{yy}D_{yy} + b_{zz}D_{zz} + (b_{xy} + b_{yx})D_{xy} + (b_{xz} + b_{zx})D_{xz} + (b_{yz} + b_{zy})D_{yz}\right)$$

Note that D is never explicitly calculated in practice, but instead is estimated based on a best fit model. This is because dMRI measurements are subject to many sources of error, including field inhomogeneities, thermal noise, motion artifacts, and eddy-current effects. Therefore, while the theoretical minimum number of diffusion scans needed to estimate a diffusion tensor is 7 (6 diffusion directions plus a b0 volume), 30 or more scans are typically acquired in practice to increase the accuracy of the estimate.

It is useful to decompose the diffusion tensor into its eigenvectors and eigenvalues (principle components). This decomposition reveals the primary direction of diffusion as the eigenvector associated with the largest eigenvalue. In a WM voxel, this may be conceptualized as the direction of the axon tract. The other two eigenvectors are orthogonal to this primary direction of diffusion, and their associated eigenvalues are the diffusion coefficients in these secondary and tertiary directions.

$$\mathbf{D} = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix} = \begin{pmatrix} v_{11} & v_{12} & v_{13} \\ v_{21} & v_{22} & v_{23} \\ v_{31} & v_{32} & v_{33} \end{pmatrix} \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} \begin{pmatrix} v_{11} & v_{21} & v_{31} \\ v_{12} & v_{22} & v_{32} \\ v_{13} & v_{23} & v_{33} \end{pmatrix}$$

λ_1 , λ_2 , and λ_3 are the three eigenvalues of the diffusion tensor $\mathbf{D}$. The three vertical columns of the matrix before the matrix of eigenvalues are the eigenvectors, and the matrix after the eigenvalues is simply the transpose of the eigenvector matrix.

The two most common metrics to characterize diffusion using DTI can be derived from the decomposed matrix: fractional anisotropy (FA) and mean diffusivity (MD). FA is an index of diffusional anisotropy in the voxel, and it ranges from 0 (indicating perfect isotropy) to 1 (indicating perfect anisotropy). MD is an index of the average diffusivity within the voxel and is simply the average of the three primary eigenvalues. Note that MD is also equivalent to the average of the three primary diffusivities before eigenvalue decomposition.

$$MD = \frac{D_{xx} + D_{yy} + D_{zz}}{3} = \frac{\lambda_1 + \lambda_2 + \lambda_3}{3}$$

$$FA = \sqrt{\frac{3}{2}} \cdot \frac{\sqrt{(\lambda_1 - MD)^2 + (\lambda_2 - MD)^2 + (\lambda_3 - MD)^2}}{\sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}}$$

Canonically, it is thought that a decrease in FA accompanied by an increase in MD indicates structural breakdown, as freer diffusion might occur when a cellular membrane is compromised. However, interpretation of FA and MD should be made with caution, as the DTI model makes several assumptions that may not be accurate in vivo. Importantly, DTI assumes Gaussian diffusion (explored in the following section) and homogenous diffusive properties within a voxel. The limitations of DTI are apparent when considering the case of crossing axon fibers within a voxel. If there are 3 crossing fibers in orthogonal directions, a diffusion tensor might estimate isotropic diffusion while the reality of diffusive properties within the voxel is far more complex. While recognizing the limitations of DTI as the simplest dMRI model, FA and MD are still informative and clinically relevant indices of neural microstructure.

2.4. Diffusion kurtosis imaging (DKI)

Diffusion kurtosis imaging (DKI) is an extension of the DTI model which removes the assumption of a Gaussian diffusion. To understand the significance of this assumption, it is helpful to consider the diffusion displacement probability distribution $P(\mathbf{r}, t)$. $P(\mathbf{r}, t)$ is a probability density function that captures the likelihood that a single molecule will diffuse a distance ($\mathbf{r}$) due to diffusion after a certain time (t). For free diffusion occurring in a uniform medium, $P(\mathbf{r}, t)$ is a Gaussian distribution which has a width dependent on the diffusion coefficient (D).

$$P(r, t) = \frac{1}{(4\pi Dt)^{3/2}} \exp\left(-\frac{r^2}{4Dt}\right)$$

This relation is assumed in the derivation of Stejskal and Tanner's equation relating signal attenuation to diffusion coefficient. Since the entire shape of the distribution is solved if D is known, only a single diffusion-weighted scan is required to estimate $P(r, t)$. However, diffusion in neural tissue is hindered by cell membranes and other biological substrate, causing $P(r, t)$ to deviate from the Gaussian. The extent of deviation from a Gaussian diffusion is known as kurtosis (K). Greater kurtosis suggests a greater degree of cell structure, as there are more barriers to free diffusion. A key consequence of having a diffusion profile with non-zero kurtosis is that a single diffusion-weighted scan is no longer sufficient to estimate D . Instead, two scans with different b -values must be obtained for a diffusion direction to quantify both D and K . dMRI datasets acquired with multiple non-zero b -values are known as multi-shell acquisitions. When fitting the observed signal in each voxel to estimate D and K , an adjustment is made to the Stejskal-Tanner equation [Jensen 2005]:

$$\ln\left(\frac{S(b)}{S_0}\right) = -bD + \frac{1}{6}b^2D^2K$$

$S(b)$ is the signal in a diffusion-weighted image of b -value equal to b , D is the diffusion coefficient, K is the diffusion kurtosis, and S_0 is the signal in a T2-weighted image ($b=0$).

It is evident from this equation that for a given diffusion direction, two b -values are necessary to solve the two variables D and K . Additionally, while kurtosis in a single diffusion direction can be quantified by a single value, a 4th order kurtosis tensor is necessary to fully describe a 3-dimensional kurtosis profile. The kurtosis tensor W_{ijkl} is a $3 \times 3 \times 3 \times 3$ matrix that describes the shape of $P(r, t)$ in each direction (x , y , and z) and how they interact with diffusion profiles of the other directions. For example, the term W_{xxxx} represents kurtosis along the x -direction, W_{xxyy} represents interactions between the kurtosis profiles of the x - and y -directions, and W_{xyyz} describes interactions between all three directions. As with the 3×3 diffusion tensor, the kurtosis tensor is symmetric (W_{xxyy} is equivalent to W_{yyxx}). This symmetry means that there are only 15 independent components out of the 81 elements of the kurtosis tensor. Thus, W_{ijkl} can be theoretically defined with 16 scans: a b_0 scan and 15 diffusion-weighted images sensitized in different directions across two b -values. As with DTI, though, much more data is acquired in practice to decrease the effect of noise and generate better estimates. Once the kurtosis tensor is estimated, it can be summarized into an average kurtosis coefficient (AKC) in the x , y , and z directions. These can be further summarized into a mean kurtosis (MK) value in each voxel by averaging AKC across the three directions. Typical values of MK in GM reported in literature range from about 0.4 - 1.2. Larger values of MK, indicating greater deviation from Gaussian diffusion, imply greater complexity of the local microstructure.

2.5. Neurite orientation dispersion and density imaging (NODDI)

A limitation of the DTI and DKI models is that they are agnostic to known neurobiological structure. In other words, they collapse complex and heterogeneous cellular microstructure into gross tensors that describe the aggregate diffusion profile of a voxel but do not have direct neurobiological meaning. In contrast, the neurite orientation dispersion and density imaging (NODDI) method is fundamentally built to model diffusion in neurobiological tissue. NODDI is a

method of quantifying the dispersion of axons and dendrites (collectively known as neurites) to offer biologically specific insight into microstructure [Zhang 2012]. The model breaks down each voxel into 3 components: the intracellular compartment, the extracellular compartment, and CSF. The total signal in each voxel is a linear combination of modeled signal from each of the 3 components, weighted by their volume fraction (how much of the voxel’s volume is taken by the component).

$$S = v_{iso}S_{iso} + v_{ic}S_{ic} + v_{ec}S_{ec}$$

Here, S is the aggregate voxel signal, S_{iso} is the CSF signal, v_{iso} is the CSF volume fraction (ISOVF), S_{ic} is the intracellular compartment signal, v_{ic} is the intracellular volume fraction (ICVF), S_{ec} is the extracellular compartment signal, and v_{ec} is the extracellular compartment volume fraction.

The intracellular compartment represents water within neurites. Their narrow, tubular structure imposes significant anisotropic restrictions on diffusion; therefore, neurites are represented as zero-radius cylinders. Conceptually, this is equivalent to a specific case of the diffusion tensor with only one nonzero eigenvalue. The NODDI model considers the intracellular signal A_{ic} to be an integral of zero-radius cylinders over a spatial distribution. The spatial distribution is defined by a spherical probability density function called the Watson Distribution (Figure 7) [Watson 1960]. A formalization of this function is outside the scope of this thesis, but importantly, it is always bidirectionally and cylindrically symmetric. These symmetries are reasonable constraints on neurite dispersion from a biological perspective, and they allow the Watson distribution’s shape to be fully defined by two variables: μ —the principal orientation direction, and κ —the degree of dispersion about μ . κ is the key parameter of NODDI: a high κ represents low dispersion and low κ represents high dispersion. The principal outcome metric of NODDI is orientation dispersion index (ODI), which is simply a transformation of κ to constrain its range to $[0, 1]$ and invert its interpretation such that 0 indicates low dispersion:

$$ODI = \frac{2}{\pi} \arctan\left(\frac{1}{\kappa}\right)$$

The extracellular compartment represents water outside of neurites in GM and WM regions. NODDI models the diffusion in this compartment as anisotropic and Gaussian with its shape determined by the Watson distribution of the intracellular compartment. The signal from the extracellular compartment is analogously an integral over the Watson distribution where anisotropic diffusion tensors with components parallel and perpendicular to the zero-radius cylinders are integrated instead of the cylinders themselves. The parallel component is equivalent to free diffusion in the direction of the neurite, while perpendicular diffusion is hindered by the neurite cell membrane. In the NODDI model, this relationship is simplified by an assumption about diffusion around cell membranes such that the two components are dependent. Therefore, a Watson distribution of neurites defined by κ fully determines the parallel and perpendicular components of extracellular diffusion tensors which are integrated across the distribution to model the extracellular signal component A_{ec} . The third and final compartment is for CSF and is simply modeled by an isotropic Gaussian diffusion profile.

To summarize, while DTI and DKI can provide hints about cellular microstructure by modeling the aggregate diffusion profile in each voxel, NODDI is built on biological priors and therefore offers indices of microstructure specific to neural tissue. In practice, NODDI requires multiple b-shells and nearly 100 total diffusion directions to accurately estimate its model parameters. ODI is the

primary outcome metric of this method, but ISOVF and ICVF (the volume fractions of CSF and intracellular components) are also commonly reported. In cortical GM, increases in ODI are associated with brain development in youth, while decreases in ODI are seen in aging elderly patients [Zhang 2012].




Figure 7: Example Watson distributions with for NODDI orientation dispersion index (ODI) values ranging from 0.04 to 1.0. NODDI uses the 3D Watson distribution to model the dispersions of axons and dendrites. Note that the distributions are cylindrically symmetrical, therefore only the 2D projections are shown. Figure reproduced from Zhang et al., 2012.

3. Modulating the cortex with TMS

While neuroimaging can non-invasively elucidate structural and functional abnormalities of the human brain in vivo, it is a purely diagnostic tool with no therapeutic effect. Other methods are needed to target and therapeutically modulate the abnormalities observed by imaging. In the context of chronic pain, several methods have been explored to varying degrees of success. The method used in the body of this thesis is a non-invasive brain stimulation (NIBS) technique known as transcranial magnetic stimulation (TMS). This section will review the advantages and disadvantages of using TMS to modulate the cortex compared to other therapies, describe the basic principles of TMS, and finally discuss the specific, recently developed implementation of TMS used in this thesis.

3.1. Therapeutic neuromodulation methods in chronic pain: why TMS?

Therapeutic modulation of the human brain in vivo can be achieved via many mechanisms. These therapies can be grouped conceptually into a few categories, including pharmacological intervention, cognitive-behavioral training, lifestyle changes, and brain stimulation therapies (invasive and non-invasive). In the context of chronic pain, common pharmacological interventions include opioids, selective serotonin reuptake inhibitors (SSRIs), non-steroid anti-inflammatory drugs (NSAIDs), and

anti-seizure medications [Kuijpers 2011, Wiffen 2013, Patetsos 2016]. While there is moderate evidence that these drugs can provide some pain relief and functional restoration, they have a global non-specific influence on the whole CNS, leading to a risk of adverse side effects. Moreover, opioids are known to be highly addictive and carry risks of overdose and drug abuse [Chou 2015]. Cognitive-behavioral interventions for chronic pain include cognitive behavioral therapy (CBT) and mindfulness meditation. These therapies focus on modulating coping strategies and psychological symptoms associated with pain, but they can also lead to mild pain relief and induce neuroplastic changes [Broderick 2016, Hilton 2017]. However, the efficacy of these interventions crucially depends on patient compliance, motivation, and expectation for the treatment, which limits its universal applicability [Broderick 2016]. Lifestyle changes that can be effective in chronic pain include exercise and diet regimens. These treatments can induce pain relief, improve mood, and beneficial neuroplastic changes [Ambrose 2015, Field 2021]. However, their implementation varies widely, and they are additionally limited by patient compliance and ability to adhere to the prescribed regimen.

Multiple brain stimulation treatments for chronic pain have been studied. Invasive stimulation therapies include deep brain stimulation (DBS), spinal cord stimulation, motor cortex stimulation, and peripheral nerve stimulation. All these methods involve surgical placement of an electrode or array of electrodes at different points in the nervous system. These electrodes are then used to directly stimulate target neurons. Stimulation of the spinal cord, peripheral nerves, and motor cortex attempt to modulate the pain signal early in its processing, while DBS can modulate downstream processing of pain by targeting areas like PAG and ACC. These invasive electrical stimulation therapies are effective at reducing pain for some patients, but a large proportion of patients do not respond to the treatment (estimated 20-50% in spinal cord stimulation, 24-33% in motor cortex stimulation, and 33% for peripheral nerve stimulation) [Boccard 2015, Helm 2021, Levy 2010, Verrills 2016]. This family of therapies also carry the risks, costs, and barriers to access associated with any surgical procedure. In contrast with invasive techniques, several NIBS therapies for chronic pain have also been explored. These include transcutaneous electrical nerve stimulation (TENS), peripheral magnetic stimulation (PMS), transcranial direct current stimulation (tDCS), transcranial alternating current stimulation (tACS), focused ultrasound stimulation (FUS), and TMS. TENS and PMS are non-invasive analogs to peripheral nerve stimulation, while tDCS and tACS modulate the excitability of the cortex by applying a current across two electrodes placed on the scalp. Current evidence surrounding tDCS and tACS therapies for chronic pain indicate inconclusive outcomes or nonsignificant effects of the interventions [Alwardat 2020, Chang 2023]. FUS is a novel technique for non-invasively stimulating deeper brain regions which shows promising effects in preclinical studies but has not been tested in clinical cohorts.

TMS utilizes the principle of magnetic induction to directly stimulate cortical neurons in a targeted area. A brief review of the basic principles of TMS is presented in the following section. Here, the advantages and disadvantages of TMS as a treatment for chronic pain are discussed in relation to the other therapies presented. TMS is non-invasive and can target a specific cortical area, eliminating complications related to invasive surgeries and globally active pharmacological treatments. It is also less dependent on patient compliance compared to cognitive-behavioral therapies and prescribed lifestyle changes, since the participant is a passive recipient of treatment in TMS rather than an active participant. Compared to non-invasive peripheral nerve stimulation therapies, TMS can modulate cortical regions involved in higher-level processing of pain which are thought to be abnormal in chronic pain states. Finally, TMS is a more clinically mature technique compared to

other cortical NIBS methods like tDCS, tACS, and FUS. Current evidence of TMS in chronic pain from clinical trials generally supports a higher degree of efficacy compared to these other techniques, although it should be noted that the efficacy of NIBS in pain is unclear even for TMS. To accurately describe how TMS fits into the space of chronic pain therapies, its disadvantages should also be noted. Notably, the specific therapeutic mechanism of TMS is still unclear, making it difficult to design and optimize treatment protocols. TMS also requires specialized equipment and knowledge to operate, making it less accessible and harder to implement compared to some other options. While the side effects of TMS are less substantial compared to pharmacological interventions and invasive stimulation procedures, there is a low chance of inducing seizure and syncope during treatment. The targeting of TMS is less focused compared to invasive brain stimulations, and it is not able to directly stimulate deeper brain targets like DBS or FUS. Considering its advantages and disadvantages holistically in the context of other chronic pain treatments, TMS occupies a unique niche in its non-invasive and targeted mechanism which holds potential for positive analgesia and functional improvements with few side effects.

3.2. Basic TMS principles

TMS is a technique for non-invasively stimulating neurons in the human cortex in vivo by leveraging laws of electromagnetic induction. Its basic principles can be demonstrated by considering the simplest TMS experiment, first described by Barker and colleagues in 1985 [Barker 1985]. In their experiment, a single pulse from a circular coil stimulator is applied over the hand knob region of the primary motor cortex (M1) to induce a muscle twitch in the contralateral hand (Figure 8). The circular stimulator is a handheld coil of wire made of a variable number of turns (ranging from less than a dozen to a few hundred) and with a typical diameter of about 5 to 10 cm. This coil is connected to a circuit with capacitors for storing electrical charge and a thyristor to quickly discharge the current through the coil stimulator. Typical TMS systems can generate thousands of amps within a few hundred milliseconds [Rotenberg 2014]. Consider a single pulse of current released by the circuit which passes through the TMS coil. According to the Biot-Savart law, this current will induce a magnetic field directed orthogonally through the center of the coil:

$$B = \frac{\mu_0 N I}{2R},$$

where B is the induced magnetic field, μ_0 is the permeability of free space (a physical constant), N is the number of turns in the wire, I is the current through the wire, and R is the radius of the wire.

Note that the strength of the induced magnetic field varies in direct proportion with the current through the coil. As the discharged current rapidly rises and falls during the pulse (dI/dt), the magnetic field also fluctuates rapidly (dB/dt), reaching about 2.5T and falling back to zero within a millisecond [Rotenberg 2014]. The magnetic pulse passes through the skull with little impedance from tissue, but its strength decreases rapidly according to the inverse cube law. Due to this rapid distance-dependent attenuation of the magnetic pulse, TMS preferentially targets superficial cortical areas. Once the magnetic pulse reaches the brain, it induces current in the neural tissue according to Faraday's law of induction.

$$\varepsilon = -N \frac{\Delta\Phi}{\Delta t},$$

where ε is the induced voltage, N is the number of turns in the wire, $\Delta\Phi$ is the change in magnetic flux, and Δt is the change in time.




Figure 8: Graphical abstract of TMS stimulation of M1. A pulse of current (dI/dt) through the TMS coil induces a magnetic pulse (dB/dt), which stimulates neurons in M1. If the TMS pulse is targeted over the hand knob region of M1, the stimulated action potentials propagate from the upper motor neuron to the lower motor neuron, inducing a muscle twitch captured by EMG. EMG: electromyography; M1: primary motor cortex; TMS: transcranial magnetic stimulation.

For a single pulse from a circular coil, the induced eddy current is circular, centered about the magnetic pulse, and parallel to the TMS coil. This current can depolarize neurons and generate action potentials which propagate to downstream pathways. TMS stimulates regions connected to the targeted area in a nonspecific manner—it is not selective for any one connection from the target region. In the case of hand knob M1 stimulation, TMS depolarizes neurons in the corticospinal tract down to the motor end plates innervating the hand muscles, generating the twitch observed in the contralateral hand. Note that M1 is somatotopically organized, with different subregions controlling distinct body parts (see Penfield's homunculus [Penfield 1937]), therefore precise targeting of the hand subregion is required to elicit the hand twitch response. Electromyography (EMG) electrodes can be placed on the abductor muscles of the target hand to record a motor-evoked potential (MEP), which reflects the aggregate activity of motor neurons stimulated by the magnetic pulse. This is the response which was recorded and reported by Barker and colleagues in their seminal 1985 experiment with single-pulse TMS over M1 (Figure 8).

Expanding upon this basic paradigm, many parameters can be modified and must be considered when designing a TMS study. In the interest of streamlining this introduction, the full explanation of TMS parameters is presented in appendix B.

3.3. Connectivity-guided intermittent theta-burst stimulation

TMS applied at a frequency of 1 Hz or higher is known as repetitive TMS (rTMS). The purported mechanism of rTMS interventions is to induce long-term potentiation (LTP) or long-term depression (LTD) of synaptic connections involving the target area. LTP refers to a sustained strengthening of a synaptic connection, while LTD refers to sustained synaptic weakening. Since the 1990's, rTMS has historically been the dominant paradigm of TMS applied in clinical trials to induce intransient neuroplasticity in human subjects. That is, rTMS is proposed to alter the strength of connections involving the target region, which may reflect LTP or LTD on the neuronal level. Recently, another specific TMS stimulation protocol known as connectivity-guided intermittent theta-burst stimulation (cg-iTBS) has been shown to have increased efficacy compared to rTMS in some clinical trials [Cole 2020]. Theta-burst stimulation (TBS) refers to TMS stimulation patterns which deliver high-frequency bursts of pulses (one burst is typically 3-5 pulses at 50-100Hz), with the interval between bursts corresponding to a frequency within the theta-band of EEG signals (4-8 Hz) [Huang 2005]. This pattern of stimulation was observed to be effective in animal studies, and indeed its high efficacy in inducing neuroplasticity was demonstrated by stimulating the human M1 and observing changes in MEP response [Huang 2005]. Intermittent TBS (iTBS) refers to theta-burst stimulation protocols which include periods of rest within the protocol which are longer than TBS stimulation (e.g. 2 seconds of stimulation followed by 8 seconds of rest). Finally, connectivity-guidance for TMS refers to the practice of using fMRI to determine the stimulation target. For example, a recent study in depression used resting-state fMRI analysis to determine the region of the IDLPFC most anticorrelated with the subgenual ACC [Cole 2020]. Then they used this individualized coordinate as a stimulation target for iTBS. The rationale behind this connectivity-guidance was to target the specific connection between IDLPFC and subgenual ACC, as modulation of this interaction was hypothesized as a possible therapeutic mechanism. Thus, combining the technique of connectivity-guidance with iTBS stimulation parameters gives a TMS protocol which is accelerated (delivers more pulses in shorter time), individualized, and reported to have greater efficacy compared to traditional rTMS designs.

3.4. TMS in chronic pain

Clinical trials have used rTMS to try to alleviate pain across a variety of conditions. A review from 2020 found 106 studies applying rTMS to 16 pain conditions, including central pain after stroke, fibromyalgia, migraine, chronic low back pain and knee OA among others [Yang 2020]. Note that these pain conditions have a wide range of etiologies and underlying mechanisms (e.g. traumatic vs. degenerative, neuropathic vs. musculoskeletal). Despite the heterogeneity of pain conditions, almost all studies have targeted one of two regions for stimulation: M1 or DLPFC. Systematic reviews have found that stimulation of M1 has weak to moderate analgesic effects [Crucchi 2016, Yang 2020]. The mechanism of analgesia induced by M1-stimulation is uncertain, but several hypotheses have been proposed. Stimulation of M1 could alter cortical excitability of the motor circuit and enhance motor cortex inhibition [Lefaucheur 2008]. The activity of further targets involved in the cortical processing of pain (e.g. anterior insula and thalamus) could also be affected by M1 stimulation [Lindsay 2021]. Finally, M1 stimulation may also act on the descending pathways projecting to the dorsal horn of the spinal cord. Note that because TMS pulses have nonspecific downstream effects and do not selectively activate specific efferent connections from the target region, all of the purported mechanisms are possible. Preliminary evidence of M1 stimulation in one patient with chronic pain in knee OA shows promising analgesic effects [Nguyen 2019]. There have been fewer

studies targeting DLPFC in chronic pain, and the efficacy of DLPFC stimulation is inconclusive [Cruccu 2016]. DLPFC stimulation is purported to relieve pain by acting on the cognitive-affective processing of pain. The DLPFC also has downstream connections to the anterior cingulate and PAG, which may affect the cognitive processing and descending modulation of pain respectively [Lindsay 2021]. While rTMS of the DLPFC has been applied in other pain conditions with inconclusive results, no studies have used this intervention in chronic pain secondary to knee OA. Furthermore, there is no published evidence to date using cg-iTBS in human pain conditions.

4. BoostCPM: A clinical trial using iTBS in patients with chronic knee OA

The body of this thesis analyzes neuroimaging data collected from the clinical trial “Brain connectivity-guided optimized theta-burst stimulation to improve central pain modulation in knee osteoarthritis (BoostCPM)” [Drabek 2023]. This is a randomized controlled trial evaluating the efficacy of a cg-iTBS protocol for modulating cortical control mechanisms in chronic pain. The cg-iTBS target is an individualized point in the IDLPFC that is most functionally anticorrelated with the rAI, defined by functional connectivity analysis of baseline functional MRI (fMRI) data. The IDLPFC-rAI connection was chosen as the target to be modulated based on evidence from a previous study showing that stimulation of this connection improves symptoms of depression [Morris 2024]. By targeting the IDLPFC, the intervention aims to modulate the cognitive-affective dimensions of pain, even if there is little effect on sensory pain. This study is novel in several ways: it is the first study using cg-iTBS in a chronic pain condition and one of the first targeting the IDLPFC with TMS in patients with knee OA.

This trial is a pilot study and therefore meant to build preliminary evidence by establishing a variance of treatment outcomes and explore mechanistic hypotheses. It should be emphasized that this study is therefore not meant to be inferentially informative about the efficacy of the cg-iTBS intervention. The mechanistic hypotheses to be explored are stated in the pre-published study protocol [Drabek 2023] and are as follows:

1. “Accelerated, personalized iTBS of the DLPFC modulates functional activity and connectivity of brain circuits involved in central pain control in chronic musculoskeletal pain compared with sham controls.”
2. “Accelerated, personalized iTBS of the DLPFC increases endogenous antinociception in chronic musculoskeletal pain compared to sham controls.”

To evaluate mechanistic hypotheses, neuroimaging data is collected at baseline and after the intervention. The neuroimaging protocol includes arterial spin labeling (ASL) and fMRI, which are used to evaluate functional neuroplasticity, and dMRI, which is used to evaluate structural neuroplasticity. This thesis analyzes the dMRI data from the BoostCPM trial.

Neuroplastic changes are hypothesized in DPMS following the cg-iTBS intervention. These regions include the IDLPFC, rAI, vmPFC, dmPFC, ACC, and PAG. A previous study using dMRI to evaluate the effects of rTMS found small microstructural changes in pertinent white matter tracts that were related to treatment efficacy [Ning 2022]. This result suggests that TMS does affect neural microstructure, and that modern dMRI techniques are sensitive enough to detect these changes. While the study by Ning et al analyzes microstructural changes in WM (as is most common in dMRI analysis), this thesis will analyze microstructure in GM regions. Indexing GM microstructure with

dMRI is typically more difficult, as the distinct anisotropy of WM regions lends to simpler analysis and interpretation. However, it has been shown that a simple learning and memory task can engender changes in the GM detectable by dMRI acquisitions [Brodt 2018]. Therefore, it is hypothesized that cg-iTBS of the lDLPFC in patients with chronic pain secondary to knee osteoarthritis induces microstructural changes in GM regions of the DPMS, detectable with dMRI. The following sections present and discuss the neuroimaging analysis evaluating this hypothesis.

II. Methodology

1. Study design

The protocol for the BoostCPM trial was pre-published [Drabek 2023]. Specifics of this protocol relevant to the present thesis are summarized in this section, including participant characteristics, the iTBS protocol, and the MRI scanning protocol.

1.1. Participants

The recruitment target for this study was 45 participants, with 30 in the active iTBS intervention arm, and 15 in the sham arm. In the end, 43 participants completed the study, with 31 in the active arm, and 12 in sham. Inclusion criteria for the study restricted the cohort to patients aged 18-75 with knee OA and chronic pain (≥ 6 month duration) with a pain score via visual analog scale of $\geq 4/10$. Participants were excluded if they had major medical or neurological conditions, imminent major changes in their treatment plan (e.g. arthroplasty), a change in pain medications within the last 4 weeks, a prescription medication acting on the central nervous system (except stable antidepressants or opioidergic analgesics), any medication increasing the risks of seizure or syncope, or experienced frequent headaches.

1.2. iTBS protocol

In total, iTBS was delivered on four separate days within a 5-day span, with one day of rest. The protocol on each of the 4 intervention days was identical, with 5 sessions of iTBS per day separated by 50 minute breaks (see Figure 9). iTBS was delivered using a neuronavigated MagStim system with a figure-eight coil (7cm diameter) with the following stimulation parameters: 3 pulses at 50Hz repeated with a inter-train interval of 200ms (5Hz) for 2 minutes, for a total of 1800 pulses at a stimulation intensity of 80% resting motor threshold (rMT). The number of pulses delivered each day across the 5 sessions was 9,000. The total number of pulses delivered across the 4 intervention days was 36,000. Each participant's rMT was determined at the start of the first visit by delivering single pulses to left M1 at increasing intensities until an MEP is evoked in the abductor pollicis brevis of the right hand in 5 out of 10 trials.

The stimulation target for iTBS was a point in the participant's IDLPFC determined by functional connectivity analysis of their baseline resting state fMRI (rs-fMRI) data. Rs-fMRI data was preprocessed, and functional connectivity was calculated using a seed region in the rAI and an anatomically masked region in the IDLPFC. The rAI seed was a 6mm spherical region of interest (ROI) centered about the Montreal Neurological Institute atlas (MNI) coordinate [30, 24, -14], taken from previous studies targeting the IDLPFC-rAI connection in depression [McGrath 2013]. The IDLPFC anatomical mask is a large regional ROI and was taken from a previous study targeting this connection for stimulation [Fox 2012]. From functional connectivity analysis, the coordinate in the IDLPFC mask most anti-correlated with the rAI seed was used as the target for stimulation. The StimGuide neuronavigation system allows this individualized stimulation point to be accurately targeted during iTBS. StimGuide is a technique for registering a participant's head position in space during stimulation to their anatomical MRI scan, thereby allowing accurate, individualized placement

of the TMS coil. In brief, cameras within the iTBS stimulation room are used to generate reference points at the participant's nasion and left/right tragus. Then, the participant's T1 image is registered to these reference points, thereby allowing for accurate, real-time guidance to of the TMS coil to the target identified with fMRI analysis.



Figure 9. Graphical overview of BoostCPM study timeline. For each recruited patient, the study takes place over 6 visits. On the first and last visits, patients receive an MRI with T1, T2-FLAIR, dMRI, and fMRI scans. During each intervention visit (days 2-5), patients receive 5 sessions of cg-iTBS with 50-minute breaks between sessions.

1.3. MRI protocol

MRI scans were acquired before and after the iTBS interventions. The scanning protocol for each time point was identical, with T1, T2-weighted-Fluid-Attenuated Inversion Recovery (T2-FLAIR), rs-fMRI, and dMRI scans being obtained. T2-FLAIR is an MRI sequence which generates a structural image with T2-weighted contrast and is designed to suppress signal from CSF. All images were acquired on a 3T GE Premier scanner with a 48-channel head coil. The two structural scans (T1 and T2-FLAIR) were acquired with 256 sagittal slices, 256x256 in-plane resolution, voxel size of 1mm isotropic, and a 3D readout. T1 images were acquired with an MP-RAGE sequence (TI=800ms, flip angle=8), and T2-FLAIR images were acquired with Cube sequence (TE=121ms,

TR=6300ms, TI=1787ms, echo train length=220). rs-fMRI scans were acquired with a gradient-echo EPI sequence: TR=1400ms, TE=35ms, flip angle=68°, multiband acceleration=3, 57 axial slices, 106x106 in-plane resolution with 2mm isotropic voxels, 643 timepoints for a total acquisition time of 15 minutes. dMRI scans were acquired with a spin-echo EPI sequence: TR=4600ms, TE=90ms, multiband acceleration=3, 117 diffusion directions, 3 shells (b=300, 1000, and 2000), 63 axial slices, 104x104 in-plane resolution with 2mm isotropic voxels. For field mapping and EPI distortion correction, images with reversed phase encoding were acquired for rs-fMRI and dMRI.

2. MRI processing and analysis

Neuroimaging data must be processed through a series of steps (known as a pipeline) before it can be analyzed. There exist many software tools and methods for performing this processing, and the best practice may vary depending on the quality of the data and the goals of the analysis. This section details the processing and analysis procedures that were applied to the neuroimaging data acquired from the BoostCPM study. These details were specified a priori in an image analysis plan which was written and uploaded to the University of Nottingham's research data repository (DOI: <http://doi.org/10.17639/nott.7388>). The relevant components of this protocol are summarized and expanded upon in this section. All image processing, quality control (QC), and analysis were performed while blinded to the subjects' treatment allocation (sham or active).

2.1. Software and setup

Image preprocessing was conducted using the BRC imaging pipeline version 1.6.3, which is openly available on Github (github.com/SPMIC-UoN/BRC_Pipeline). The pipeline implements tools from various imaging software packages, including Statistical Parametric Mapping 12 (SPM12), FMRIB Software Library (FSL), Advanced Normalization Tools (ANTs), and Freesurfer. EDDY motion correction and NODDI analysis use the NVIDIA CUDA Toolkit. Processing and analysis were run on a combination of servers (Chilli) and high-performance computing clusters (Augusta and Ada) at the University of Nottingham.

Raw data is output from the MRI scanner in the DICOM data format. These data files were converted to NIFTI files using the dcm2niix software, as the NIFTI data format is more compact and has greater compatibility with image processing software. Then, the NIFTI files were organized and renamed according to the brain imaging data structure (BIDS) specification [Gorgolewski 2016]. BIDS is a standard data format created by consensus from the neuroimaging research community. Its purpose is to standardize file naming and data organization formats to facilitate comprehension when sharing data with other researchers. The preprocessing subsequently described is performed on the raw, BIDSified NIFTI data.

2.2. Structural processing and QC

T1 and T2-FLAIR data were processed using the *struc_preproc.sh* script within the BRC pipeline. There are six steps in the T1 processing pipeline: cropping to center the brain, brain extraction, nonlinear registration to the MNI152 template, bias-field correction, segmentation into major tissue

types (GM, WM, and CSF), and cortical parcellation. A graphical summary of structural processing is shown in Figure 10.




Figure 10. Graphical abstract of structural processing. T1 images are cropped, brain extracted, non-linearly registered to MNI space, bias-field corrected, tissue segmented, and parcellated with Freesurfer. T2-FLAIR images are cropped, brain extracted, and linearly registered to T1 space, then the downstream results from T1 processing are applied to the T2 image in T1-space.

1. **Cropping the image** to center the brain removes extraneous voxels below the level of the brainstem, thereby reducing the file size for future computation. This step is necessary because the field of view (FOV) of raw structural images includes the entire head down to approximately the C5 level. Cropping was performed using FSL's *robust_fov*.
2. **Brain extraction** isolates the neural tissue in the image, which helps facilitate later processing steps. Brain extraction was performed using FSL's brain extraction tool (*BET*). This step generates a binary mask of the brain and removes the skull and rest of the head from the data. *BET* was run with the “-R” flag, which recursively estimates the brain mask to improve robustness. The “-g” and “-f” flags (which modulate the brain extraction to be more liberal or conservative in certain directions) were left to their default values. Alternate options for this step include the *3dSkullStrip* tool in the Analysis of Functional NeuroImages (AFNI) toolbox and *antsBrainExtraction* in the ANTs toolbox. Note that in the BRC pipeline, while *BET* is used to isolate the brain prior to nonlinear registration, the brain mask is updated after registration. This is done by applying the inverse warp field (MNI-to-T1) to the brain mask in MNI space. The additional step is performed because brain mask estimations from this process are generally more robust than the raw estimates from *BET*.
3. **Nonlinear registration** warps the subject's structural scan to match an anatomical template. The most common template for adult neuroimaging studies is the MNI152 template from the Montreal Neurological Institute based on an average of 152 scans of healthy adults taken from the International Consortium for Brain Mapping (ICBM) project [Mazziotta 2001, Fonov 2009]. Registration to a template is necessary to facilitate comparisons between subjects, as it provides a standard space to compare individuals with different neuroanatomy. There are two general types of registration: linear and nonlinear. Linear registrations are defined by a 4x4 affine matrix which specifies translations, rotations, scalings, and shears in the x, y, and z directions. This type of registration is commonly used for intra-subject registrations (e.g. aligning T1 scans at two time points), as its parameters are sufficient to align brains with the same anatomy. However, when registering brains of different individuals (i.e. inter-subject registration), nonlinear registration is necessary. Nonlinear registrations are defined by warp fields and are therefore more flexible to deform anatomy. Registration to the MNI152 template requires nonlinear registration because the template is an averaged brain with blurred anatomy that is distinct from every individual. Two methods of nonlinear registration to the MNI152 template were performed. The first method is FMRIB's Nonlinear Image Registration Tool (*FNIRT*) from the FSL toolbox. This algorithm calculates a warp field from an input image to a target image, using cubic splines to fit the transformation and a sum of squared differences cost function to optimize the registration [Andersson 2007]. The second method is the *antsRegistrationSyn* tool in the ANTs toolbox. This algorithm uses a symmetric diffeomorphic transformation, which differs from the *FNIRT* method in two key ways [Avants 2006]. First, it considers the bidirectional registrations equally, optimizing for both the registration from source to target image and the reverse registration from target to source. Second, it ensures the calculated warp field is diffeomorphic, meaning that both the field and its inverse (the reverse transformation) is

smooth and differentiable. Note that the ANTs method was shown to generate more accurate registrations than *FNIRT* in a study comparing 12 different registration algorithms [Ou 2014], but both methods were tested in the present analysis to validate the approach. Prior to running *FNIRT*, FSL’s Linear Image Registration Tool (*FLIRT*) was used to estimate an affine matrix for the T1-to-MNI transformation. This affine matrix was used to initialize *FNIRT*, a step which accounts for parts of the transformation and therefore decreases the required intensities of the warp field. The interpolation (resampling) method used for *FNIRT* was spline. For the ANTs approach, *antsRegistrationSyn* was implemented with the “-d” flag set to 3, indicating 3-dimensional registration (as opposed to in-plane).

4. **Bias-field correction** adjusts the intensities of voxels in the image to account for the bias field, which is estimated empirically during this step. The bias field refers to the field inhomogeneities that are present in every MRI scanner. Recall from the discussion of MRI physics (Introduction section 2.1.) that all protons in the scanner are assumed to be exposed to the same net magnetic field at rest. In practice, the field inside the scanner will vary due to hardware limitations. Additionally, when subjects are placed in the scanner, further inhomogeneities are introduced due to magnetic susceptibility effects. These inhomogeneities systematically modulate voxel intensities and must be corrected for in preprocessing. Bias-field correction was performed by the FMRI’s Automated Segmentation Tool (*FAST*) in the FSL toolbox [Zhang 2001]. This method assumes that there are 3 tissue classes in the image (GM, WM, and CSF), which each have a different Gaussian voxel intensity distribution with distinct means and a small width [Guillemaud 1997]. The algorithm then iteratively updates an estimated bias field to adjust intensities in each voxel to best fit the assumed distribution of intensities. Thus, the bias-field correction method implemented in FSL *FAST* is purely data-driven (only requiring a T1-image as an input) and also segments the image into GM, WM, and CSF voxels.
5. **Segmentation into major tissue** types creates masks for GM, WM, and CSF. These masks are useful in later processing steps to isolate specific tissue types. This step was performed using FSL *FAST* using the method described for bias-field correction above.
6. **Cortical parcellation** further segments the cortical GM into distinct regions defined by topographical or functional characteristics. The method used for cortical parcellation is the Freesurfer software [Fischl 2012]. Freesurfer contains a comprehensive 29-step pipeline that includes preprocessing of structural data, fits a surface-based topological model, and outputs multiple segmentations. For the purposes of cortical parcellation, it should be noted that the segmentations that Freesurfer produces are based primarily on topology by a procedure which divides a surface representation of the cortex into sulci and gyri [Desitieux 2010]. Freesurfer was run with the *recon-all* command on the raw, unprocessed T1 images, subjecting the raw data to the entire pipeline of Freesurfer processing and modeling.

T2-FLAIR data were processed following the T1 data. The processing of T2-FLAIR data was largely analogous to that of T1, and many of the results from T1 processing were simply applied to the T2-FLAIR data. There were 5 steps in the T2 processing procedure: cropping to center the brain, brain extraction, linear registration to T1-space, application of T1-to-MNI nonlinear registration results, and bias field correction. Cropping and brain extraction were performed on the raw T2-FLAIR data as described above for the T1 data. Linear registration to T1-space was performed using *FLIRT*. Note that linear registration is used for this step because the brain anatomy is identical across the two images. Instead of recomputing the nonlinear registration from T2-to-MNI, the warp field estimated from T1 processing is simply applied to the T2-FLAIR data. Similarly, the bias field estimated from T1 processing is applied to the T2-FLAIR data.

QC of structural preprocessing was performed qualitatively, with focus on visually evaluating the accuracies of brain extraction, MNI registration, and tissue segmentation. For each dataset, FSLEYES was used to overlay the brain extraction mask on the processed T1 image, assessing for discrepancies (e.g. an overly liberal or conservative brain mask). FSLEYES was also used to inspect GM, WM, and CSF masks generated by FSL FAST. Finally, MNI registration was evaluated by using FSL's *slivesdir* tool to overlay GM-WM boundaries of the registered T1 image with the MNI atlas.

2.3. dMRI processing and QC

While several of the basic preprocessing steps performed in structural processing must also be applied to dMRI data (e.g. brain extraction and spatial registration), there are several unique challenges presented by diffusion-weighted acquisitions which require specialized tools. These challenges arise from two aspects of the dMRI acquisition: the fast EPI readout and the rapidly switching gradient coils. Recall that dMRI uses fast EPI readouts to enable collection of more gradient directions within a given amount of time and reduce motion artifacts, and that this method is susceptible to geometric distortions along the phase-encoding direction due to magnetic field inhomogeneities (see Introduction section 2.2). In addition, each diffusion-weighted scan requires the application of two pulsed gradients in a specified direction. The constant on-off switching on of the gradient coils leads to induced electric currents in the scanner hardware which can introduce distortions into the image. Therefore, specific methods for correcting these sources of distortion are necessary in the dMRI processing pipeline. Diffusion data were preprocessed using the *dMRI_preproc.sh* script within the BRC pipeline. There are 5 steps in the dMRI preprocessing pipeline: normalization of b0 intensities, distortion correction based on field map estimation, brain extraction, motion correction and eddy current distortion correction, and registration to T1- and MNI-space. A graphical summary of dMRI processing is shown in Figure 11.




Figure 11. Graphical abstract of dMRI processing. The raw dMRI dataset (117 volumes across 3 non-zero b-values) has b0 volume intensities normalized, is distortion-corrected using the field map estimation method FSL *topup*, brain-extracted, motion and eddy current distortion-corrected using FSL *eddy*, linearly registered to T1 space and nonlinearly registered to MNI space, and fitted with the DTI, DKI, and NODDI models. The methods of FSL *topup* and FSL *eddy* are shown to the right. FSL *topup* uses dMRI images acquired with reverse phase encoding to estimate a field map, correcting for gross geometric distortions due to field inhomogeneity. FSL *eddy* uses an iterative method of estimating eddy current and motion parameters, using a Gaussian predictor to evaluate the accuracy of the estimates during each epoch.

1. **Normalization of b0 intensities** rescales the intensities of all b0 volumes within a dMRI acquisition to be consistent. In a dMRI dataset, multiple b0 volumes are acquired for the purpose of improving signal-to-noise ratio (SNR) by averaging across scans. For the data collected in this study, 7 b0 volumes are collected for the dMRI scan. The global average of these scans may drift due to slight variations in scanning conditions, for example from hardware performance or thermal noise. By normalizing the intensity of each b0 scan to its average and rescaling to the mean intensity of the first b0 volume, the effect of this global drift is diminished. This step is performed using various FSL tools to perform basic mathematical operations.
2. **Distortion correction based on field map estimation** is a method of correcting EPI geometric distortions by using two images with opposite phase-encoding to estimate magnetic field inhomogeneities [Andersson 2003]. dMRI data in the Boost study is acquired with anterior-posterior (A/P) phase-encoding, which compresses the image according to deviations in the field along the A/P axis. The second, brief acquisition of dMRI data uses the same parameters but performs phase-encoding in the P/A direction, thereby stretching the image in the other direction along the same axis. The FSL tool *topup* uses information from these two scans to estimate the underlying field map which causes these distortions (see Figure 11 for a graphical representation). It uses an iterative method which begins by guessing a field map, applying the distortion correction to both images, calculating the difference between the two images, then making a new prediction based on the result to minimize the difference. After many iterations, the algorithm converges on an estimated field map which most closely aligns the two images when correcting distortions using the map. FSL *topup* is run using b0 volumes from the A/P and P/A dMRI scans. The output is a distortion-corrected b0 volume and the estimated field map.
3. **Brain extraction** is performed on the distortion-corrected b0 volume from *topup* in preparation for registration to T1-space. This step is performed with FSL *BET*, analogous to structural processing.
4. **Motion correction and eddy current distortion correction** account for two important sources of noise in dMRI (subject motion and gradient-induced eddy currents) by modeling their effects on the acquired data. They are separate processes but can be simultaneously modeled and corrected using the method implemented in the FSL tool *eddy* [Andersson 2016]. The crux of this method is a prediction algorithm that models the diffusion profile (and corresponding dMRI signal intensities) in each voxel as a Gaussian process and adjusts the model to fit observed data while being constrained by a parametrized covariance function [Andersson 2015]. This dMRI prediction maker is leveraged in the context of distortion correction preprocessing by using it in an iterative algorithm to update guesses for eddy current and motion parameters. The iterative algorithm begins with processing each diffusion-weighted image in the dMRI dataset by applying the field map from *topup* and a guess of the eddy current and motion parameters. The “distortion-corrected” images are then fed into the prediction maker, which estimates the diffusion profile in each voxel. Note that the predicted data will differ from the input images because they reflect the estimated

diffusion profile generated by the Gaussian process, which is encouraged to be smooth and axially symmetric by the imposed covariance function. The inverse of the guessed distortion correction is then applied to the predicted images, and a difference map is calculated between the original data and the result is compared to the initial data. This difference map is used to update the estimated eddy current and motion distortion terms. As the estimated terms get closer to their true values, the difference between the “distortion-corrected” (with guessed parameters) images and the images output by the prediction maker will become more similar. Thus, the difference maps will tend towards 0 as the process converges on an optimal set of distortion correction parameters. In practice, *eddy* was applied with a few notable options, including: the “--repol” flag to replace outlier volumes with their Gaussian process predictions, the “--flm=quadratic” flag to model eddy-induced currents as quadratic, and the “--mporder=4” flag to perform slice-to-volume correction (used for multiband acquisitions).

5. **Registration to T1- and MNI-space** was performed to enable cross-subject comparisons and utilization of segmentation from T1-processing. Diffusion-to-T1 registration was linear and was performed with FSL *FLIRT*. The source volume for dMRI-space was the averaged b0 data, and the target volume in T1-space was the processed T1 data. This registration used a boundary-based registration cost function, which uses the WM segmentation in T1-space to optimize alignment [Greve 2009]. After linear registration to T1-space was performed, the nonlinear registration warp field calculated during structural preprocessing was applied to the dMRI data to transform it to the standard MNI152 space.

After preprocessing, the *dMRI_preproc.sh* script subsequently modeling of the data. Diffusion tensors were fitted according to the DTI and DKI models described in Introduction sections 2.3 and 2.4. DTI modeling was performed using the FSL implementation *dtifit*, while DKI modeling used the *dipy* (Diffusion Imaging in Python) toolbox. NODDI modeling was performed with the CUDA Diffusion Modeling Toolbox (*cuDIMOT*).

QC of dMRI data was performed quantitatively and qualitatively. Quantitative QC metrics were generated with the FSL QUAD software as part of *eddy* preprocessing [Bastiani 2019]. For each dataset, QUAD calculates absolute motion, average relative motion (including decomposition into between and within-volume translations and rotations), number of outlier volumes, linear terms of eddy current distortions, SNR of b0 volumes, and contrast-to-noise ratio (CNR) of all other volumes. Then, the participant-level QUAD reports are combined into a group-level SQUAD report showing the distributions of each metric across the cohort. The group-level distributions are also split between pre-*i*TBS and post-*i*TBS sessions to determine if any systematic variation in data quality exists between sessions. Pre-post comparisons were also performed intra-subject. Datasets with outlier QC metric values within the group distribution, defined as greater than two standard deviations from the mean, were flagged for visual inspection. Finally, the group-level distributions were compared to distributions from the UK Biobank study as a benchmark for overall study-level image quality.

Qualitative QC of dMRI data focused on assessing for issues with distortion correction, FOV truncations, and motion correction. Raw and processed dMRI data were overlaid and visually

inspected with FSLEYES, with particular focus on the frontal polar region of the brain where susceptibility effects are most prominent due to proximity to the air-filled paranasal sinuses. Any datasets with issues (e.g. imperfect distortion correction, truncated brain regions) were flagged but only removed from subsequent analysis if the issue was seen to affect the extraction of final metrics.

2.4. ROI definitions

After processing structural and dMRI data, ROIs must be defined to specify the parts of the brain that will be analyzed with statistical testing. In practical terms, an ROI is a binary image data file with dimensions matching the target space. For example, an ROI in T1-space will have the same dimensions as the T1 image with voxels of value “1” indicating points within the region of interest and voxels of value “0” indicating all points outside the region. As specified by the mechanistic hypotheses stated in Introduction section 4, the brain regions to be analyzed will be PAG, ldlPFC, rAI, bilateral rACC, bilateral dmPFC, and bilateral vmPFC. These are key components of the DPMS which may be structurally modulated by the cg-iTBS intervention. In addition to these 11 anatomically defined ROIs, the GM stimulation target and two control regions are included for analysis. All ROIs were defined in T1-space, and multiple techniques for defining the binary ROIs for these regions were tested and validated (Figure 12). As an experimental analysis, ROIs were also transformed to a subject-specific T1-midspace, but this additional step was ultimately removed to minimize registration-associated errors.

Two atlases were tested for defining cortical ROIs: the multimodal Glasser atlas and the surface-based Freesurfer parcellation. The Glasser atlas is based on architectural, functional, connectivity-based, and topographic information and contains 180 cortical parcellations per hemisphere [Glasser 2016]. While the primary implementation of this atlas is as a machine learning classifier which outputs a parcellation given structural data and resting-state fMRI, it has also been translated to a volumetric atlas in MNI space. The volumetric MNI version of the Glasser atlas was used for analysis (see Discussion section 3.1. for explanation). The Freesurfer atlas is based solely on topographical features and contains 74 parcellations per hemisphere [Destrieux 2010]. Freesurfer uses structural information from anatomical MRI to generate subject-specific parcellations in T1-space. Note that there are two Freesurfer atlases: the Desikan-Killiany atlas from 2006 and the Destrieux atlas from 2010 [Desikan 2006, Destrieux 2010]. The latter is used in this thesis due to its finer parcellations (74 labels per hemisphere vs. 35 in the Desikan-Killiany atlas).

Using the Glasser approach, parcels of the MNI atlas corresponding to each of the cortical ROIs were identified based on a combination of literature and identification by a neuroradiologist (DPA). These ROIs were then nonlinearly transformed to each subject’s T1-space by using FSL *applywarp* to apply the nonlinear warp field calculated by ANTs during structural preprocessing. The masks in T1-space are probabilistic (since the transformation requires interpolation) and therefore must be binarized. Multiple thresholds for binarization were tested, and a cutoff of 75% was chosen based on empirical evaluation of the resulting masks. Using the Freesurfer approach, parcels corresponding to each of the cortical ROIs were identified by a neuroradiologist (DPA). Since the parcellation is made in native T1-space, nonlinear transformation is not necessary in this approach. However, the Freesurfer parcellations of IDLPFC, dmPFC, and vmPFC are broad and encompass areas outside those relevant to pain processing. To narrow the scope of these ROIs, they are

intersected with 1cm radius spherical ROIs centered on relevant coordinates taken from the pain neuroimaging literature. For dmPFC, the MNI coordinates [-12, 50, 34] and [12, 50, 34] were used [Zhang 2021]. For vmPFC, the MNI coordinates [-6, 45, -6] and [6, 42, -12] were used [Yu 2014]. The iTBS stimulation target ROI was defined by intersecting a cortical GM mask from Freesurfer with a 1cm sphere centered on the stimulation coordinate calculated from functional connectivity analysis. Note that the stimulation target was only calculated for 17/28 of the participants in the active group due to logistical issues and for 0/12 of the participants in the sham group.




Figure 12. ROI processing steps for the Glasser and Freesurfer approaches. Top: for the Glasser approach, all ROIs are initially defined in MNI space and must be transformed into T1 space by applying the nonlinear warp field calculated during structural processing. Then, the ROIs are thresholded based on an empirical threshold of 0.75. Bottom: for the Freesurfer approach, only the dmPFC and vmPFC ROIs need to be processed, as the others are simply defined in T1 space. The extent of the dmPFC and vmPFC Freesurfer ROIs (yellow) is limited by intersecting them with a 1cm radius spherical ROI centered on coordinates taken from literature (blue) to generate the final focal ROI (red).

For the subcortical PAG, the chosen ROI was taken from a published atlas which probabilistically mapped subcortical structures using 7T MRI [Keuken 2014]. In contrast with binary ROIs, voxels in probabilistic ROIs may have values ranging from [0, 1] indicating the likelihood that that voxel belongs to the region (where 1 indicates 100% chance of belonging). This atlas included probabilistic ROIs for young, middle-aged, and elderly populations to account for systematic age-related variations in anatomy. The “elderly” ROI was chosen, as this age range most closely matches the demographics for this clinical trial. Since this ROI was defined in the standard MNI space, it was nonlinearly transformed to each subject’s T1-space. The transformed ROI in T1-space was then thresholded at a probability of 100% (cutoff determined empirically) to binarize the probabilistic map.




Figure 13. Final ROIs in T1 space. *Left:* mPFC and rACC ROIs using the Freesurfer and Glasser approaches. Note that the Freesurfer image in this column also includes the region of no interest (cuneus, yellow). The Glasser ROIs from left to right in this image are: anterior_24_prime_L (white), area_posterior_24 (light blue), area_24 (orange), area_9_Middle (beige), area_10r (pink), area_10v (red). The Freesurfer ROIs from left to right in this image are: mid-anterior ACC (red), anterior ACC (white), dmPFC (dark blue), and vmPFC (light blue). *Middle:* rAI ROIs using the Freesurfer and Glasser approaches. *Right:* PAG ROI, defined from an independent atlas without Freesurfer or Glasser.

Two control ROIs included in the analysis are a nonspecific cortical GM mask and the cuneus. The nonspecific cortical GM mask includes all cortical GM and serves to check for global drifts in the dMRI signal, which are not expected. However, this mask includes regions of interest where hypothesized changes could bias the global GM signal. Therefore, the cuneus (part of the occipital lobe) was chosen as a region of no interest. The cuneus is anatomically and functionally distinct from the other 11 regions of interest included above. It is also not expected to be involved in the processing or modulation of pain signals. The ROI for the cuneus was defined from 5 Freesurfer regions which were concatenated to form a large region encompassing the entire lobe.

A visualization of all ROIs generated by the Glasser and Freesurfer methods is shown in Figure 13.

2.5. Extraction of dMRI outcomes

Following the analysis of preprocessed dMRI data using the DTI, DKI, and NODDI models, relevant imaging metrics were extracted from the ROIs defined. The extraction was performed using a custom script written in Python3. To perform the extraction, binary ROI masks were used to index the MD, MK, ODI, ISOVF, and ICVF maps in T1-space, generating a distribution of values corresponding to the relevant region. For the MD distribution, values outside the range of biologically plausible values (0.0005-0.001 mm²/sec) were excluded [Brodt 2018]. The histogram of values was plotted for each distribution of metrics in each ROI for each subject, and these were visually inspected to assess for abnormalities. Large, implausible shifts or skewed distributions suggest a possible methodological issue which warrants visual inspection. If the distributions do not suggest any methodological issues, the median of the distribution is extracted. The median is chosen as the outcome statistic because it is less susceptible to being skewed by outliers than the mean. Thus, after extraction of the median value of 5 dMRI metrics (MD, MK, ODI, ISOVF, and ICVF) from 14 ROIs (PAG, rAI, cuneus, cortical GM, IDLPFC, iTBS target, bilateral anterior ACC, bilateral mid-ACC, bilateral dmPFC, and bilateral vmPFC), each subject will have 70 median diffusion statistics which can be compared pre- and post-iTBS.

3. Statistical testing

All statistical tests are performed on an exploratory basis and are not inferential, as the BoostCPM study was designed as a pilot without prior knowledge of variance in each metric of interest. Furthermore, tests on the 12 predefined regions of interest (the iTBS targeted connection, cortical DPMS regions, and PAG) are interpreted separately from the tests on the two control regions (nonspecific cortical GM and cuneus) as they serve a different purpose. Tests on the regions of interest were predefined in the registered image analysis plan and are hypothesized to show an effect. In contrast, tests on the two control regions constitute a sensitivity analysis to compare against preplanned test. The control regions are not expected to show an effect. Note that the frequentist statistical framework used does not allow for acceptance of the null hypothesis. Nonsignificant results are interpreted as absence of evidence for an effect, not as evidence for no effect. This distinction is important for the multiple comparison correction method explained later in this section.

With these caveats in mind, statistical testing evaluated the significance of two contrasts per metric: a within-group pre-post comparison and a between-group difference-of-differences comparison. The within-group comparison considers the active and sham groups separately and tests whether there is a difference in each dMRI metric between the pre-*i*TBS and post-*i*TBS time points. For example, one of these tests will evaluate whether there is a difference in MD in the rAI from pre- to post-*i*TBS for patients in the active group. In total, 135 of these tests are performed: 14 ROIs x 5 dMRI metrics x 2 groups, without the *i*TBS target ROI in sham. Formally, these tests are two-tailed paired t-tests, with hypotheses:

$$H_0: x_{\text{post}} - x_{\text{pre}} = 0$$

$$H_1: x_{\text{post}} - x_{\text{pre}} \neq 0$$

where x_{post} is the median dMRI metric in the ROI post-*i*TBS, and x_{pre} is the median dMRI metric in the ROI pre-*i*TBS.

The between-group comparison tests whether the magnitude of pre-post change is different between the active and sham conditions. In total, 65 of these tests are performed: 13 ROIs x 5 dMRI metrics (again, the *i*TBS target is not included because there is no sham data). Formally, these tests are two-tailed independent samples t-tests, with hypotheses:

$$H_0: \Delta x_{\text{active}} - \Delta x_{\text{sham}} = 0$$

$$H_1: \Delta x_{\text{active}} - \Delta x_{\text{sham}} \neq 0$$

where Δx_{active} is the pre-post difference in the dMRI metric within the specified ROI for the active group ($x_{\text{post, active}} - x_{\text{pre, active}}$), and Δx_{sham} is the pre-post difference in the dMRI metric within the specified ROI for the sham group ($x_{\text{post, sham}} - x_{\text{pre, sham}}$).

All significance tests will initially be conducted with a significance threshold of $p=0.05$, defined a priori in the image analysis plan. After the initial round of testing, p-values for the 12 regions of interest will be corrected for multiple comparisons. Although each t-test specified involves distinct sets of data and no distribution is tested against multiple contrasts, the high number of tests performed means that the probability of observing a false positive in this suite of tests is inflated. To counteract this inflation, p-values will be adjusted using Benjamini and Hochberg's method to adjust the false discovery rate (FDR) to 0.05 [Benjamini 1995]. That is, any significant result surviving the correction will have a 5% chance of being a false positive. Note that this is a more liberal approach to the multiple comparisons problem compared to other common methods such as Bonferroni or Sidak corrections, which broadly adjust for the probability of observing any false positives in the set of tests. Due to the relatively small expected effect size, sample size, and number of tests (compared to, say, voxelwise fMRI designs) in this study, the comparatively liberal FDR method is used here. FDR correction is only performed for t-tests involving the 12 regions of interest. T-tests for the two control regions are not included in the correction because they serve as a qualitative benchmark to compare the other results against, and their results are interpreted in an entirely distinct way.

III. Results

1. Quality control of MRI data and analysis

1.1. Data exclusion and harmonization

Out of the 43 participants who completed the study, 41 were deemed to have a complete and usable set of MRI data. Two datasets (subjects 002 and 006) were excluded because their pre-iTBS dMRI images were acquired with different parameters compared to all other data in the study. The pre-iTBS dMRI data for these datasets were acquired with a much higher in-plane resolution than other datasets (256x256 instead of 106x106 for all other data). To ensure data homogeneity throughout the cohort and prevent pre-post comparisons between datasets acquired with different parameters, these subjects were removed from analysis. One other dataset (subject 035) was excluded during preprocessing after encountering a dataset-specific software error during dMRI processing. Thus, data from 40/43 participants who completed the study were fully processed and used for analysis.

1.2. QC of structural preprocessing

QC of structural data revealed slight errors in BET, tissue segmentation, and nonlinear registration to MNI space, but none of the errors were deemed to significantly affect downstream results (Figure 14). Checking the accuracy of BET by overlaying the extracted brain mask over the T1 image revealed small truncations in the most superior part of the parietal lobe. These truncations were on the order of a few voxels and affected 16/40 datasets. Slight truncations can affect downstream processing in two ways: if we are extracting outcome metrics from the truncated area or if it alters spatial registration. This region of the parietal lobe was not defined as a region of interest in the analysis plan since it is not expected to be active in pain. Therefore, these slight truncations may only affect spatial registration, which is also checked in the QC process.

Registration with both ANTs and FNIRT was checked by overlaying the cortical GM boundary with the T1 image transformed to MNI space. When considering the GM boundary, registrations from both ANTs and FNIRT were observed to deviate from MNI in various datasets. Note that registration is not expected to always result in perfectly aligned GM boundaries because this will require large distortions in some datasets. In other words, registration algorithms should not overfit data and eliminate all true variations in anatomy. However, FNIRT resulted in registrations that warped the brain into anatomically unreasonable constructions in 5/40 datasets (see Figure 14, bottom). Furthermore, ANTs resulted in better-aligned cortical GM in most of the remaining 35 datasets. Therefore, ANTs was chosen as the registration method for this pipeline.

Tissue segmentation with FAST was checked by overlaying the GM, WM, and CSF masks on the T1 data. Cortical GM was observed to be segmented well in all datasets, but subcortical GM was not captured by the algorithm. This is an expected result, as FAST is known to have poor results in subcortical GM due to the low contrast with surrounding WM voxel intensities. Poor segmentation in subcortical GM is not expected to affect downstream processing, as the tissue segmentations are only used as an initialization for linear registration of dMRI to T1 using boundary-based registration.




Figure 14. Quality control of structural processing in three steps. *Top left:* Brain extraction (BET) is checked by overlaying the BET mask (red) over the T1 image. Very slight truncations in the cortex (green arrows) are noted in 16/40 datasets, but these are not expected to affect downstream results. *Top right:* Tissue segmentation of grey matter (GM), white matter (WM), and cerebrospinal fluid (CSF) is checked by overlaying the GM mask (red) and CSF mask (dark blue) over the T1 image. Tissue segmentation with FSL *FAST* is observed to miss subcortical GM structures (green crosshair), but this is not expected to affect downstream results. *Bottom:* Nonlinear registration to MNI space is checked by overlaying the GM boundary of the MNI template (red traces) over the transformed T1 images in MNI space. For this subject (sub-016), the T1 image registered by ANTs has cortical folding that deviates from MNI (yellow arrows). However, the T1 image registered by FNIRT has anatomically unreasonable constructions (blue arrows).

1.3. QC of dMRI preprocessing

Quantitative QC of dMRI data using the FSL QUAD software (part of *eddy*) generated group-level distributions of SNR for b0 volumes, CNR for all b-shells, absolute motion and relative motion parameters, and percentage of volumes classified as outliers. The group-level distributions for the BoostCPM study were compared to distributions from a subset of the UK Biobank study generated by Bastiani et al [Bastiani 2019] (Figure 15, top). The UK Biobank was chosen for comparing the distributions because its acquisition parameters are similar to those used in the BoostCPM study (both studies used 2mm isotropic voxel size, b1000 and b2000 shells with 50 diffusion directions each). The SNR and CNR of dMRI scans from the present study are slightly lower but within reasonable ranges. The estimated relative motion parameters are lower than the UK Biobank distribution, indicating less head movement on average. The distributions of QC metrics were also compared between all pre-iTBS scans and post-iTBS scans to assess for bias across time points (Figure 15, middle). No difference in the mean and variance of these distributions was apparent. However, a few datasets with large motion stood out from the others—these were more carefully inspected. The time series of motion estimates, raw data, and processed data were visualized (Figure 15, bottom). Volumes associated with large movements were often seen to have large artifacts in the raw data, however these are not present in the processed data as *eddy* replaces these volumes with predicted data from the fitted Gaussian process (see Figure 11, bottom right). Thus, it was determined that even datasets with large spikes in motion were usable for subsequent analysis.

Qualitative QC of dMRI data using FSLEYES revealed distortion-correction errors in the frontal polar region of the brain (Figure 16). The errors could be classified as either blurring of the distortion-correction region (32/40 datasets) or undercompensation of the spatial unwarping (6/40 datasets). Blurring was characterized by spatial unwarping that appeared anatomically faithful, but with signal intensities in the undistorted region that dropped out and appeared smoothed (Figure 16, top). This was present in most datasets and is an expected feature of the *topup* distortion correction method in regions of high inhomogeneity. However, it is notable that this blurring propagated to the calculated MD map. The affected region is anterior to the vmPFC ROI defined by Freesurfer but near enough to warrant cautious interpretation of the results from this area. Undercompensation of spatial unwarping was present in only a few datasets, but these represent a more serious failure of *topup*. The estimated field map in these datasets was insufficient to fully unwarped the spatial distortions, leading to an anatomically compromised image (Figure 16, bottom). It is apparent from the distortion-corrected image that unwarping has not been successful, as evidenced by persistent hyperintensities that indicate aggregation of dMRI signal. As with the blurring errors, undercompensation of spatial unwarping also propagates to errors in the MD map in the frontal polar region.

QC with FSLEYES also revealed FOV truncations of the superior parietal lobe of the brain for 26/40 datasets, akin to those seen in the QC of structural processing (e.g. Figure 14, top left). This region is not used for extraction of dMRI data; therefore, the truncations should not affect the results. There were no issues with between-volume motion correction observed.




Figure 15. Quantitative quality control of dMRI processing with FSL *eddy*'s quality assessment tool (QUAD). QUAD outputs average SNR for b0 volumes, average CNR for all $b > 0$ shells, and estimated motion parameters for each dataset, then plots the group-level distribution of these metrics. **Top:** group-level dMRI QC metrics are compared to a subset of the UK Biobank. While SNR and CNR is slightly lower for the Boost study, motion is generally lower in the Boost cohort (except a few outliers). **Middle:** group-level dMRI QC metrics are compared for all pre-iTBS vs. post-iTBS scans. No systematic differences in the distributions are seen, but a few scans are noted as outliers. Subject 001, ses-1 is marked with a red dot. **Bottom:** outlier datasets are inspected by examining the evolution of their motion parameters and visually inspecting volumes with large translations. While a large artifact is seen in the raw data, *eddy* replaces this scan with the Gaussian prediction.




Figure 16. Qualitative assessment of distortion correction during dMRI processing reveals blurry frontal lobes where high levels of distortions exist in the raw image, and undercompensation on some datasets. These errors are seen to carry through to the calculated MD maps and are limited to the frontal polar region.

1.4. Comparison of Freesurfer and Glasser ROIs

While the ROIs defined by Freesurfer were seen to closely follow the cortical GM ribbon, the Glasser ROIs transformed to T1 space deviated significantly from anatomical boundaries (Figure 17). It was observed during QC of structural processing that although ANTs was a superior method to FNIRT for nonlinear registration to MNI space, it still generated transformations with misaligned cortical GM ribbons. Therefore, the reverse transformation (MNI-to-T1) applied to the Glasser ROIs in MNI space were subject to the same misalignment. Because the diffusion characteristics of GM, WM, and CSF are vastly different, it is imperative that chosen ROIs closely follow the GM boundary. For this reason, Freesurfer ROIs were chosen for final extraction of dMRI metrics.




Figure 17. Glasser ROIs transformed into T1 space do not align with the cortical GM ribbon. The 6 Glasser ROIs representing rACC, dmPFC, and vmPFC are all observed to have poor alignment with the cortical GM. Yellow lines indicate the plane corresponding to the images on the right.

2. Structural neuroplasticity

The outcome indices of iTBS-induced structural neuroplasticity are MD from the DTI model, MK from the DKI model, and ODI, ICVF, and ISOVF from the NODDI model. Changes in these metrics are evaluated by model for each of the 14 ROIs in the analysis. The ROIs are split into four conceptual groupings: two control regions (cortical GM and cuneus), three regions involving the targeted IDLPFC-rAI connection (IDLPFC, iTBS target, and rAI), eight regions involving the cortical DPMS (bilateral dmPFC, vmPFC, anterior ACC, and mid-anterior ACC), and finally the subcortical PAG.

2.1. Neuroplastic changes: MD

The results for changes in median MD values across the active and sham treatment groups are summarized in Figure 18. Results of statistical tests from within-group comparisons are presented in Table 1.1., and results from between-group comparisons are presented in Table 1.2.

At the subject level, there is high variability in the direction and magnitude of pre-post differences observed for both the active and sham groups in all 14 ROIs (Figure 18, panel C for active group, panel D for sham group). Median MD is seen to increase in some subjects and decrease in others—no region shows a consistent direction of effect. Some subjects demonstrate large changes in median MD extracted from vmPFC regions which exceed 0.0001 mm²/sec. This magnitude of change is physiologically unlikely and may be suggestive of errors attributable to methodological variability, especially considering that these regions are closest to areas of distortion-correction errors (see Figure 18).

At the group level, distributions of MD values for the pre- and post-iTBS time points (Figure 18, panel A for active group, panel B for sham group) are largely similar in most ROIs, with some regions showing slight changes in the median (e.g. cuneus, iTBS target, rAI, lvmPFC). However, the range and inter-quartile range of the iTBS target and rAI distributions are largely similar between the two time points. There is a starker difference in the overall distributions for lvmPFC. Note that the variance in distributions ranges widely—it is tighter in the large cortical GM ROI and looser in the vmPFC ROIs closest to the high-distortion frontal polar area. There is an observable difference between the median MD distribution in cortical GM and cuneus, suggesting that the cuneus may not be an appropriately defined control region.

Statistical testing indicated within-group significance in the active group for lvmPFC and cortical GM (Figure 18, panel E). In both cases, there was an increase in MD from pre- to post-iTBS. The significance in lvmPFC does not survive FDR correction (recall that FDR correction is not performed for control regions for conceptual reasons, see methods section 3). The MD increase in cortical GM appears to be driven by increases in IDLPFC, rAI, vmPFC, and ACC regions. Pre-post differences are again most variable in vmPFC regions compared to other ROIs. Note how the confidence intervals for the sham group are generally larger due to the smaller sample size. While none of the other groups reached statistical significance, the confidence intervals for the active group seem to skew more positive compared to the sham group, which generally seem to be more centered about $\Delta=0$.




Figure 18. MD results for 14 ROIs including control regions, the iTBS-targeted connection, cortical DPMS, and PAG. **A)** Boxplot distributions for each of the 14 ROIs, active group only. Pre-iTBS median MD distributions are in blue, and post-iTBS distributions are in beige. **B)** Same as plot A, but for the sham group. **C)** Raw data points for each participant in the active group, pre-post change indicated by lines connecting points from the same subject. **D)** Same as plot C, but for the sham group. **E)** Magnitude of pre-post changes for the active (green) and sham (red) groups. Individual pre-post changes are plotted as points. The shaded areas represent the 95% confidence interval of pre-post change from statistical testing. Significant results ($p < 0.05$ before multiple comparison correction) for within-group comparisons are indicated by a star.

Table 1.1: Within-group statistical testing results for MD (Post - Pre)

ROI	Group	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	Active	4.7E-07	7.2E-06	0.03*	0.62
Cortical GM	Sham	-2.7E-06	7.8E-06	0.31	0.83
Cuneus	Active	-2.8E-07	8.7E-06	0.07	0.62
Cuneus	Sham	-6.8E-06	1.0E-05	0.66	0.92
IDLpFC	Active	-5.3E-06	1.7E-05	0.30	0.83
IDLpFC	Sham	-1.2E-05	1.5E-05	0.84	0.97
iTBS target	Active	-8.1E-06	1.9E-05	0.40	0.87
iTBS target	Sham	N/A	N/A	N/A	N/A
rAI	Active	-5.8E-06	1.2E-05	0.48	0.90
rAI	Sham	-1.6E-05	2.4E-05	0.68	0.92
ldmPFC	Active	-1.1E-05	1.3E-05	0.89	0.97
ldmPFC	Sham	-1.3E-05	1.9E-05	0.68	0.92
rdmPFC	Active	-1.0E-05	1.3E-05	0.80	0.95
rdmPFC	Sham	-8.4E-06	2.6E-05	0.29	0.83
lvmPFC	Active	4.0E-06	3.0E-05	0.01*	0.62
lvmPFC	Sham	-2.3E-05	2.2E-05	0.97	0.99
rvmpFC	Active	-2.0E-06	3.3E-05	0.08	0.62
rvmpFC	Sham	-4.2E-05	2.8E-05	0.68	0.92
lAntACC	Active	-5.4E-06	1.4E-05	0.39	0.87
lAntACC	Sham	-1.1E-05	2.4E-05	0.43	0.87
rAntACC	Active	-5.1E-06	1.5E-05	0.32	0.83
rAntACC	Sham	-1.2E-05	2.1E-05	0.57	0.92
lMidAntACC	Active	-4.2E-06	5.9E-06	0.73	0.94
lMidAntACC	Sham	-4.6E-06	1.2E-05	0.33	0.83
rMidAntACC	Active	-8.4E-06	8.1E-06	0.97	0.99
rMidAntACC	Sham	-4.3E-06	7.3E-06	0.58	0.92
PAG	Active	-5.0E-06	8.3E-06	0.62	0.92
PAG	Sham	-7.5E-06	1.7E-05	0.41	0.87

Table 1.2: Between-group statistical testing results for Δ MD (Active - Sham)

ROI	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	-4.8E-06	7.3E-06	0.66	0.92
Cuneus	-6.9E-06	1.2E-05	0.59	0.92
lDLPFC	-1.2E-05	2.1E-05	0.60	0.92
rAI	-2.2E-05	2.1E-05	0.94	0.95
lDmPFC	-2.2E-05	1.7E-05	0.80	0.95
rdmPFC	-2.7E-05	1.3E-05	0.46	0.92
lvmPFC	-7.7E-06	4.3E-05	0.16	0.92
rvmPFC	-1.6E-05	6.1E-05	0.23	0.92
lAntACC	-2.1E-05	1.7E-05	0.80	0.95
rAntACC	-1.8E-05	1.9E-05	0.94	0.95
lMidAntACC	-1.3E-05	6.5E-06	0.51	0.92
rMidAntACC	-1.1E-05	8.1E-06	0.73	0.92
PAG	-1.7E-05	1.0E-05	0.64	0.92

2.2. Neuroplastic changes: MK

The results for changes in median MK values across the active and sham treatment groups are summarized in Figure 19. Results of statistical tests from within-group comparisons are presented in Table 2.1., and results from between-group comparisons are presented in Table 2.2.

At the subject level, there is low consistency in the direction and magnitude of pre-post change, similar to what was observed for MD (Figure 19, panels C and D). In contrast to the raw MD values, the relative magnitude of pre-post change is higher for many of the ROIs except the control regions. Pre-post changes in the vmPFC ROIs appear especially large for MK as well. In one subject, the MK values for the iTBS target ROI appears to be an outlier and may warrant more careful methodological inspection.

At the group level, distributions of MK values for the pre- and post-iTBS time points (Figure 19, panel A for active group, panel B for sham group) are again largely similar in most ROIs. The variability in MK median distributions is larger across different ROIs than across the two time points. The median MK value appears elevated from pre- to post-iTBS for the rAI in the sham group. There is again a difference in the median MK distribution in cortical GM and cuneus, further suggesting that the cuneus ROI may not be representative of unaffected cortical GM.

Statistical testing indicated a within-group significance in the active group for the left mid-anterior ACC (Figure 19, panel E). The direction of change was an increase in MK from pre- to post-iTBS. The p-value did not survive FDR correction. Large confidence intervals are again observed for the bilateral vmPFC ROIs. Some of the confidence intervals for the sham group appear to skew positive while the active group for the same ROI is centered about $\Delta=0$ (e.g. in rAI, bilateral vmPFC, and bilateral anterior ACC).




Figure 19. MK results for 14 ROIs including control regions, the iTBS-targeted connection, cortical DPMS, and PAG. **A)** Boxplot distributions for each of the 14 ROIs, active group only. Pre-iTBS median MK distributions are in blue, and post-iTBS distributions are in beige. **B)** Same as plot A, but for the sham group. **C)** Raw data points for each participant in the active group, pre-post change indicated by lines connecting points from the same subject. **D)** Same as plot C, but for the sham group. **E)** Magnitude of pre-post changes for the active (green) and sham (red) groups. Individual pre-post changes are plotted as points. The shaded areas represent the 95% confidence interval of pre-post change from statistical testing. Significant results ($p < 0.05$ before multiple comparison correction) for within-group comparisons are indicated by a star.

Table 2.1: Within-group statistical testing results for MK (Post - Pre)

ROI	Group	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	Active	-0.003	0.016	0.16	0.77
Cortical GM	Sham	-0.016	0.018	0.91	0.98
Cuneus	Active	-0.009	0.015	0.64	0.92
Cuneus	Sham	-0.036	0.009	0.22	0.77
IDLpFC	Active	-0.025	0.038	0.68	0.92
IDLpFC	Sham	-0.025	0.045	0.54	0.92
iTBS target	Active	-0.016	0.047	0.32	0.83
iTBS target	Sham	N/A	N/A	N/A	N/A
rAI	Active	-0.047	0.039	0.85	0.97
rAI	Sham	-0.003	0.089	0.06	0.62
ldmPFC	Active	-0.046	0.028	0.61	0.92
ldmPFC	Sham	-0.066	0.051	0.78	0.95
rdmPFC	Active	-0.020	0.054	0.36	0.84
rdmPFC	Sham	-0.058	0.051	0.89	0.97
lvmPFC	Active	-0.067	0.050	0.77	0.95
lvmPFC	Sham	-0.047	0.101	0.43	0.87
rvmpFC	Active	-0.051	0.072	0.73	0.94
rvmpFC	Sham	-0.073	0.156	0.45	0.87
lAntACC	Active	-0.038	0.034	0.91	0.98
lAntACC	Sham	-0.037	0.073	0.49	0.90
rAntACC	Active	-0.034	0.033	0.98	0.99
rAntACC	Sham	-0.017	0.079	0.19	0.77
lMidAntACC	Active	0.005	0.049	0.02*	0.62
lMidAntACC	Sham	-0.008	0.040	0.17	0.77
rMidAntACC	Active	-0.001	0.037	0.06	0.62
rMidAntACC	Sham	-0.018	0.040	0.43	0.87
PAG	Active	-0.012	0.064	0.17	0.77
PAG	Sham	-0.010	0.044	0.19	0.77

Table 2.2: Between-group statistical testing results for Δ MK (Active - Sham)

ROI	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	-0.013	0.024	0.54	0.92
Cuneus	-0.009	0.041	0.19	0.92
lDLPFC	-0.049	0.042	0.88	0.95
rAI	-0.107	0.013	0.12	0.92
ldmPFC	-0.069	0.065	0.95	0.95
rdmPFC	-0.043	0.083	0.52	0.92
lvmPFC	-0.127	0.055	0.42	0.92
rvmPFC	-0.156	0.095	0.62	0.92
lAntACC	-0.083	0.043	0.52	0.92
rAntACC	-0.088	0.025	0.27	0.92
lMidAntACC	-0.021	0.043	0.49	0.92
rMidAntACC	-0.026	0.041	0.65	0.92
PAG	-0.036	0.054	0.69	0.92

2.3. Neuroplastic changes: NODDI

The results for changes in median ODI, ICVF, and ISOVF values across the active and sham groups are summarized in Figures 20, 21, and 22, respectively. Results of statistical tests from within-group comparisons are presented in Table 3.1. (ODI), Table 4.1. (ICVF), and Table 5.1. (ISOVF). Statistical results from between-group comparisons are presented in Table 3.2. (ODI), Table 4.2. (ICVF), and Table 5.2. (ISOVF).

At the subject level, pre-post changes in ODI appear to be relatively higher in the iTBS target ROI compared to other regions. The magnitude of change in the vmPFC regions does not appear to be significantly larger than in the other ROIs for ODI. This is in contrast with the data for ICVF, where there are exceptionally high magnitudes of change in the vmPFC regions. Considering the raw ISOVF data, there appear to be more outlier datasets which deviate from the other data for the given ROI.

At the group level, the distribution of ODI values in PAG is starkly lower than for all other regions. This may be due to the PAG having a simpler structure with strong preference for superior-inferior neurite direction since it sits in the brainstem, in contrast with all 13 other ROIs which are in cortical GM. The distributions for ODI appear to vary more across ROIs than across the two time points. There is little variation in ICVF distributions either across ROIs or across time points. The distribution of ISOVF values in cuneus is starkly higher than for all other regions. This suggests that the cuneus ROI may be contaminated by CSF, likely due to the proximity of this region to the cerebral aqueduct. There are some differences in the median of the distributions for ISOVF in both the sham and active groups, but also note that the variance in these distributions is quite large.

Statistical testing indicated between-group significance in ODI for the cortical GM and cuneus. The direction of change for both tests indicates greater pre-post change in ODI in the active group compared to the sham group. Note that these contrasts are generated from bidirectional effects: a slight increase in ODI for the active group and a slight decrease in ODI for the sham group are observed, neither of which are significant within-group. Statistical testing also indicated within-group significance in ICVF for the cortical GM active group. The direction of change for this result was a slight decrease in ICVF. Note that this effect is very small, and significance is aided by the low variance observed. FDR correction is not performed on any of the three significant NODDI results because they all occurred in control regions. It is worth pointing out that the confidence intervals for ODI in rAI and PAG are skewed heavily positive for the active group while being centered about $\Delta=0$ for the sham group, even though testing did not reach significance. Also note the larger confidence intervals for vmPFC regions for ICVF and ISOVF.




Figure 20. ODI results for 14 ROIs including control regions, the iTBS-targeted connection, cortical DPMS, and PAG. **A)** Boxplot distributions for each of the 14 ROIs, active group only. Pre-iTBS median ODI distributions are in blue, and post-iTBS distributions are in beige. **B)** Same as plot A, but for the sham group. **C)** Raw data points for each participant in the active group, pre-post change indicated by lines connecting points from the same subject. **D)** Same as plot C, but for the sham group. **E)** Magnitude of pre-post changes for the active (green) and sham (red) groups. Individual pre-post changes are plotted as points. The shaded areas represent the 95% confidence interval of pre-post change from statistical testing. Significant results ($p < 0.05$ before multiple comparison correction) for between-group comparisons are indicated by a star with brackets.

Table 3.1: Within-group statistical testing results for ODI (Post - Pre)

ROI	Group	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	Active	-0.000	0.004	0.09	0.62
Cortical GM	Sham	-0.010	0.001	0.08	0.62
Cuneus	Active	-0.000	0.009	0.06	0.62
Cuneus	Sham	-0.018	0.0001	0.05	0.62
IDLpFC	Active	-0.009	0.006	0.66	0.92
IDLpFC	Sham	-0.011	0.010	0.95	0.99
iTBS target	Active	-0.016	0.015	0.95	0.99
iTBS target	Sham	N/A	N/A	N/A	N/A
rAI	Active	-0.002	0.017	0.10	0.62
rAI	Sham	-0.028	0.016	0.59	0.92
ldmPFC	Active	-0.009	0.009	0.98	0.99
ldmPFC	Sham	-0.020	0.013	0.69	0.92
rdmPFC	Active	-0.003	0.013	0.24	0.77
rdmPFC	Sham	-0.014	0.016	0.88	0.97
lvmPFC	Active	-0.009	0.012	0.78	0.95
lvmPFC	Sham	-0.024	0.015	0.60	0.92
rvmpFC	Active	-0.022	0.005	0.20	0.77
rvmpFC	Sham	-0.027	0.002	0.09	0.62
lAntACC	Active	-0.009	0.006	0.69	0.92
lAntACC	Sham	-0.013	0.009	0.66	0.92
rAntACC	Active	-0.011	0.005	0.45	0.87
rAntACC	Sham	-0.020	0.003	0.12	0.69
lMidAntACC	Active	-0.005	0.007	0.70	0.92
lMidAntACC	Sham	-0.015	0.010	0.63	0.92
rMidAntACC	Active	-0.007	0.007	0.99	0.99
rMidAntACC	Sham	-0.022	0.002	0.09	0.62
PAG	Active	-0.002	0.022	0.09	0.62
PAG	Sham	-0.040	0.010	0.21	0.77

Table 3.2: Between-group statistical testing results for Δ ODI (Active - Sham)

ROI	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	0.001	0.013	0.02*	0.79
Cuneus	0.003	0.024	0.01*	0.74
lDLPFC	-0.013	0.011	0.83	0.95
rAI	-0.010	0.037	0.25	0.92
lDmPFC	-0.015	0.021	0.72	0.92
rDmPFC	-0.012	0.020	0.63	0.92
lvmPFC	-0.015	0.027	0.55	0.92
rvmPFC	-0.015	0.022	0.71	0.92
lAntACC	-0.012	0.014	0.90	0.95
rAntACC	-0.008	0.019	0.39	0.92
lMidAntACC	-0.009	0.017	0.54	0.92
rMidAntACC	-0.003	0.023	0.13	0.92
PAG	-0.001	0.052	0.06	0.92




Figure 21. ICVF results for 14 ROIs including control regions, the iTBS-targeted connection, cortical DPMS, and PAG. **A)** Boxplot distributions for each of the 14 ROIs, active group only. Pre-iTBS median ICVF distributions are in blue, and post-iTBS distributions are in beige. **B)** Same as plot A, but for the sham group. **C)** Raw data points for each participant in the active group, pre-post change indicated by lines connecting points from the same subject. **D)** Same as plot C, but for the sham group. **E)** Magnitude of pre-post changes for the active (green) and sham (red) groups. Individual pre-post changes are plotted as points. The shaded areas represent the 95% confidence interval of pre-post change from statistical testing. Significant results ($p < 0.05$ before multiple comparison correction) for within-group comparisons are indicated by a star.

Table 4.1: Within-group statistical testing results for ICVF (Post - Pre)

ROI	Group	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	Active	-0.008	-0.000	0.04*	0.62
Cortical GM	Sham	-0.015	0.002	0.14	0.74
Cuneus	Active	-0.005	0.010	0.52	0.92
Cuneus	Sham	-0.033	0.010	0.27	0.82
IDLpFC	Active	-0.012	0.010	0.82	0.97
IDLpFC	Sham	-0.028	0.017	0.58	0.92
ITBS target	Active	-0.013	0.016	0.82	0.97
ITBS target	Sham	N/A	N/A	N/A	N/A
rAI	Active	-0.020	0.007	0.36	0.84
rAI	Sham	-0.034	0.008	0.21	0.77
ldmPFC	Active	-0.014	0.012	0.88	0.97
ldmPFC	Sham	-0.028	0.013	0.44	0.87
rdmPFC	Active	-0.009	0.016	0.61	0.92
rdmPFC	Sham	-0.033	0.005	0.14	0.75
lvmPFC	Active	-0.045	0.004	0.10	0.62
lvmPFC	Sham	-0.029	0.007	0.21	0.77
rvmpFC	Active	-0.060	0.005	0.09	0.62
rvmpFC	Sham	-0.032	0.050	0.64	0.92
lAntACC	Active	-0.022	0.002	0.11	0.64
lAntACC	Sham	-0.031	0.007	0.20	0.77
rAntACC	Active	-0.021	0.002	0.10	0.62
rAntACC	Sham	-0.028	0.007	0.22	0.77
lMidAntACC	Active	-0.006	0.008	0.77	0.95
lMidAntACC	Sham	-0.019	0.001	0.09	0.62
rMidAntACC	Active	-0.010	0.007	0.75	0.95
rMidAntACC	Sham	-0.012	0.003	0.24	0.77
PAG	Active	-0.006	0.011	0.54	0.92
PAG	Sham	-0.026	0.014	0.49	0.90

Table 4.2: Between-group statistical testing results for ΔICVF (Active - Sham)

ROI	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	-0.007	0.011	0.60	0.92
Cuneus	-0.009	0.037	0.21	0.92
lDLPFC	-0.020	0.029	0.69	0.92
rAI	-0.018	0.031	0.57	0.92
lDmPFC	-0.017	0.030	0.57	0.92
rDmPFC	-0.005	0.039	0.13	0.92
lvmPFC	-0.039	0.020	0.51	0.92
rvmPFC	-0.087	0.014	0.15	0.92
lAntACC	-0.020	0.024	0.86	0.95
rAntACC	-0.019	0.021	0.93	0.95
lMidAntACC	-0.002	0.021	0.10	0.92
rMidAntACC	-0.008	0.014	0.59	0.92
PAG	-0.012	0.030	0.38	0.92




Figure 22. ISOVF results for 14 ROIs including control regions, the iTBS-targeted connection, cortical DPMS, and PAG. **A)** Boxplot distributions for each of the 14 ROIs, active group only. Pre-iTBS median ICVF distributions are in blue, and post-iTBS distributions are in beige. **B)** Same as plot A, but for the sham group. **C)** Raw data points for each participant in the active group, pre-post change indicated by lines connecting points from the same subject. **D)** Same as plot C, but for the sham group. **E)** Magnitude of pre-post changes for the active (green) and sham (red) groups. Individual pre-post changes are plotted as points. The shaded areas represent the 95% confidence interval of pre-post change from statistical testing.

Table 5.1: Within-group statistical testing results for ISOVF (Post - Pre)

ROI	Group	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	Active	-0.004	0.007	0.58	0.92
Cortical GM	Sham	-0.011	0.009	0.87	0.97
Cuneus	Active	-0.001	0.022	0.06	0.62
Cuneus	Sham	-0.048	0.013	0.23	0.77
IDL PFC	Active	-0.011	0.012	0.92	0.98
IDL PFC	Sham	-0.010	0.027	0.33	0.83
iTBS target	Active	-0.008	0.018	0.42	0.87
iTBS target	Sham	N/A	N/A	N/A	N/A
rAI	Active	-0.026	0.011	0.42	0.87
rAI	Sham	-0.010	0.026	0.35	0.84
ldmPFC	Active	-0.026	0.013	0.48	0.90
ldmPFC	Sham	-0.033	0.024	0.74	0.95
rdmPFC	Active	-0.019	0.016	0.86	0.97
rdmPFC	Sham	-0.023	0.016	0.68	0.92
lvmPFC	Active	-0.046	0.012	0.25	0.78
lvmPFC	Sham	-0.045	0.035	0.78	0.95
rvmpFC	Active	-0.049	0.018	0.36	0.84
rvmpFC	Sham	-0.036	0.041	0.89	0.97
lAntACC	Active	-0.027	0.009	0.32	0.83
lAntACC	Sham	-0.017	0.015	0.94	0.99
rAntACC	Active	-0.020	0.007	0.31	0.83
rAntACC	Sham	-0.016	0.025	0.61	0.92
lMidAntACC	Active	-0.004	0.019	0.17	0.77
lMidAntACC	Sham	-0.026	0.034	0.78	0.95
rMidAntACC	Active	-0.006	0.013	0.46	0.89
rMidAntACC	Sham	-0.021	0.013	0.64	0.92
PAG	Active	-0.000	0.038	0.05	0.62
PAG	Sham	-0.015	0.041	0.33	0.83

Table 5.2: Between-group statistical testing results for Δ ISOVF (Active - Sham)

ROI	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	-0.009	0.013	0.66	0.92
Cuneus	-0.004	0.060	0.08	0.92
lDLPFC	-0.029	0.013	0.44	0.92
rAI	-0.040	0.010	0.22	0.92
lDmPFC	-0.035	0.031	0.89	0.95
rDmPFC	-0.023	0.027	0.86	0.95
lvmPFC	-0.060	0.036	0.62	0.92
rvmPFC	-0.067	0.031	0.46	0.92
lAntACC	-0.032	0.015	0.46	0.92
rAntACC	-0.035	0.012	0.31	0.92
lMidAntACC	-0.028	0.035	0.81	0.95
rMidAntACC	-0.012	0.026	0.44	0.92
PAG	-0.027	0.038	0.72	0.92

IV. Discussion and Conclusion

These results do not suggest that iTBS of the DLPFC has a detectable influence on the microstructure of DPMS GM regions in patients with chronic pain in knee OA. None of the 135 statistical tests performed were close to reaching the significance threshold of 0.05 after multiple comparisons correction. Four of the six tests which reached significance before correction were in control regions, contradicting the hypothesis of microstructural change specific to the DPMS.

It is prudent to be cautious in interpreting results from dMRI data in patient cohorts due to the inherent indirectness of diffusion outcome metrics and the many layers of data processing which can all add variance to the final quantitative outputs. This is especially true with a relatively small sample size as in the present dataset. Note that even as data processing is meant to filter unwanted noise and isolate the underlying signal of interest, flexibility in processing methods means that different conclusions can be drawn from the same dataset when using distinct pipelines with individually justifiable methods [Poldrack 2017, Veraart 2022]. Nevertheless, the insights into neural microstructure offered by modern dMRI methods may have great utility for informing the development of new therapies like iTBS if analyses are properly scrutinized. This discussion will therefore begin by considering the robustness of the results obtained, then explore possible interpretations and clinical implications, suggest alternative methods and future analyses, and finally offer concluding remarks.

1. Robustness of results

1.1. Pipeline variability and QC considerations

Within the anatomical preprocessing pipeline, the only significant step where variations could be introduced is in the nonlinear registration to MNI space. Both FNIRT and ANTs were explored for this step. ANTs produced more accurate results out of these two methods, but it may be possible to generate a better registration either with a different software package or with different parameters within ANTs. The impact of variability in nonlinear registration would have been greater if ROIs were defined in MNI space and then transformed to structural space, utilizing the warp field calculated during registration. Instead, using Freesurfer ROIs defined in T1 space eliminates this extra transformation. Therefore, the effect of variability in nonlinear registration is negligible for the present analysis.

However, using Freesurfer constrained the definition of regions to those defined by this software, which is primarily based on topography. There are multiple ways to segment the brain (e.g. based on topography, histology, functional activity, or connectivity), and other atlases draw regional boundaries at different locations. For example, the location and extent of the rAI varies significantly between the Freesurfer parcellation, Glasser atlas, and Harvard-Oxford cortical atlas (Figure 17). The choice of atlas used to define ROIs can greatly affect results. It should be emphasized that the reason for using Freesurfer in this analysis was primarily to reduce variability attributable to nonlinear registration errors and not a preference for its segmentation methodology.

The processing of dMRI data followed a standard pipeline within the FSL software. Other software options exist for processing dMRI data which may lead to slightly different results in the outcome maps. Considering group-level QC, the distribution of quantitative QC metrics suggests that the raw

data from this study reaches field standards (Figure 15). However, qualitative QC revealed inadequacies in the distortion correction step which propagated to final maps (Figure 16). These errors primarily affected the mPFC region; therefore, interpretation of outcomes from the vmPFC and dmPFC ROIs should be done with some caution.

In summary, the key takeaways from considering pipeline variability and QC results are:

1. Using Freesurfer ROIs eliminated the effect of variability in nonlinear registration.
2. ROI boundaries can vary across different atlases, and this variability can influence results.
3. The interpretation of dMRI metrics in the mPFC may be less robust due to issues with distortion correction.

These caveats should be kept in mind for interpretation of all results.

1.2. Validity of quantitative dMRI results

The MD values obtained lie within a biologically plausible range: 0.0005 - 0.001 mm²/sec. However, MK values in this study drifted below the lower bound of what is typically reported for GM: 0.4 - 1.2 is commonly published, while certain ROIs in this study exhibited MK < 0.4. Lower MK values could potentially be explained by partial voluming of the ROIs with CSF, as CSF would be expected to exhibit lower kurtosis (more Gaussian-like diffusion) compared to GM. The ISOVF data lends some support to this partial voluming hypothesis. Some outlier data points exhibit abnormally high ISOVF values, suggesting increased presence of CSF within the ROI. The ODI values for cortical GM regions lie within a plausible range (0.4 - 0.6 in this dataset). The PAG is the only subcortical GM ROI, and it has notably lower ODI than the cortical GM regions. This may reflect the relatively simpler organization of the brainstem compared to cortex. Low SNR in the brainstem compared to cortical regions may also contributed to the lower ODI values in PAG. Although there is greater concern for CSF partial voluming for the PAG ROI due to its proximity to the cerebral aqueduct, ISOVF values for the PAG are not notably higher than the other regions. Together, the reasonably lower ODI values and non-elevated ISOVF values for the PAG increase confidence in this region's quantitative outcomes.

There is notably high variability in all 5 metrics within the vmPFC and dmPFC regions. This observation is consistent with the errors in distortion correction around this region noted during QC. Taken together, the high variability and distortion correction errors decrease confidence in results from the mPFC regions. It is therefore difficult to discern whether the significant increase of MD in lvmPFC before multiple comparisons correction reflects true neuroplasticity or is an artifact of poor data. Future studies are advised carefully check the integrity of dMRI data from this area.

2. Interpretation and implications

2.1. Interpretation of quantitative results

While none of the six tests which reached statistical significance were significant following correction for multiple tests, the shifts suggested by these tests may still be suggestive and update posterior beliefs about neuroplasticity in iTBS. The tests resulting in pre-correction p-values below the 0.05 threshold are presented in Table 6.

Table 6: Statistical tests with pre-correction significance

ROI	Metric	Comparison	95% CI, Low	95% CI, High	p-value	Adjusted p-value
Cortical GM	MD	Within-group, Active	4.7E-07	7.2E-06	0.03*	0.62
lvmPFC	MD	Within-group, Active	4.0E-06	3.0E-05	0.01*	0.62
lMidAntACC	MK	Within-group, Active	0.005	0.049	0.02*	0.62
Cortical GM	ODI	Between-group	0.001	0.013	0.02*	0.79
Cuneus	ODI	Between-group	0.003	0.024	0.01*	0.74
Cortical GM	ICVF	Within-group, Active	-0.008	-0.000	0.04*	0.62

Before speculating upon possible interpretations, though, a few points are worth raising. It should be emphasized again that BoostCPM was a pilot study, and that these statistical tests are meant to be exploratory and non-inferential. In addition, note that all the significant tests within-group occurred in the active group, which has over two times as many subjects as the sham group. It is therefore possible that the two groups have the same underlying effect size, but that the active group simply has enough subjects to generate a significant result while the sham group does not. Furthermore, two observations about these six results tend to suggest that they are spurious. First, four of the six significant results occur in control regions (cortical GM and cuneus) which are rather nonspecific regions and not expected to be affected by the intervention. These are also large ROIs with lower variance of medians, therefore smaller changes may be detected as significant. Second, significant within-group results are not corroborated by test of other metrics in the same region. An increase in MD might be expected to be accompanied by a decrease in MK or an increase in ODI, all suggesting increased structural disorganization. Instead, the significant results stand alone, increasing suspicion that they are an artifact of having 135 t-tests. Finally, a cursory look at the p-values adjusted for multiple tests also suggests that the results are spurious. With these cautions in mind, one may entertain some interpretations of the results.

The two significant tests for MD were within-group increases following active iTBS in the cortical GM and lvmPFC. An increase in MD suggests greater diffusivity within the ROI, canonically indicating a less structured architecture. This result conflicts with the hypothesized reduction in MD. The significant result within nonspecific cortical GM for this metric is interesting for a few reasons. First, cortical GM was a control region not expected to change from pre- to post-iTBS. Second, this ROI encompasses by far the greatest number of voxels, diminishing the variance in its medians. Within the ROIs evaluated, increases in MD (mostly not significant) observed in ldlPFC, rAI, vmPFC, and ACC appear to contribute to this result. It is also possible that iTBS influences cortical GM microstructure in regions outside the predefined DPMS ROIs, and future work may look at which regions of cortical GM are driving this result. It is also possible that this is a spurious result arising from random chance or unknown, systematic errors in methodology. The increase of MD in

lvmPFC should be interpreted with caution, given that this ROI is part of the mPFC region possibly affected by distortion correction issues. If the result is believed, the increased MD could signify structural breakdown within the region, potentially due to reorganization induced by iTBS. The widespread change in MD across cortical GM could indicate a large nonspecific effect of iTBS on GM structure facilitated by second, third, or higher degree connections to distant cortical regions. Given that TMS stimulates widespread connections and a large number of pulses are delivered in the iTBS protocol, neuronal reorganization may be occurring even at distant sites. In addition to cortical GM, increase in MD is also seen in the lvmPFC. The lvmPFC is part of the mPFC region which can influence PAG and the DPMS. Therefore, true microstructural changes in this region would be an important result, especially if accompanied by patient-reported analgesia.

The significant test for MK was a within-group increase in the lMidAntACC ROI following active iTBS. An increase in MK suggests a more structured architecture. The anterior ACC is a key node for cognitive pain processing, and microstructural change in this region could suggest a change in the salience and affective processing of nociception. Since low MK values suggest structural breakdowns associated with cognitive aging, it could be construed that an increase in MK is a positive marker of healthy function within the lMidAntACC region, although this interpretation is rather speculative. As with all the other “significant” dMRI results, the increase in lMidAntACC MK would be more interesting and important if it is correlated with patient-reported improvements.

The two significant results for ODI were between-group comparisons suggesting a greater post-iTBS increase in the active arm compared to the sham arm for both control regions: cortical GM and cuneus. Again, these were large, nonspecific areas not expected to be affected by the iTBS intervention. The within-group comparisons were not significant. An increase in ODI indicates greater dispersion of neurites, possibly suggesting neural reorganization. It is difficult to construe a mechanism by which iTBS of the IDLPFC would increase ODI in the cuneus, which is not only spatially distant but also functionally distinct from the stimulated region. This increases suspicion in a spurious result in the cuneus, and since the cortical GM ROI encompasses the cuneus, it is possible that the other significant result is driven by the observation in the cuneus. Note that significance in the cuneus is more likely to affect the cortical GM ROI than the other ROIs because it is a much larger region.

The significant result for ICVF was a within-group decrease in the cortical GM following active iTBS. This result seems to suggest that less volume in the cortical GM was taken up by intracellular space following iTBS. It is difficult to propose a mechanism by which this might occur. Also note that this result is not accompanied by a similarly significant increase in ISOVF. The lack of a plausible mechanism and isolation of this result suggest that it is spurious.

2.2. Clinical implications

While the two pre-adjusted significant p-values in lvmPFC and lMidAntACC would be interesting and important in the context of chronic pain, it is difficult to trust these results given the limitations discussed in the previous section. It is most reasonable to simply conclude that iTBS of the IDLPFC does not produce a measurable change in the microstructure of DPMS GM regions. However, the lack of a positive result from this analysis does not rule out the possibility of other effects on neural microstructure. It is possible that the effect is localized in a different region of the brain, that our outcome metrics (MD, MK, ODI, ISOVF, and ICVF) measure the wrong indices of neural structure, or that our techniques are not sensitive enough to capture the effect. Future analyses on

this dataset may be warranted to explore potential effects of iTBS on different regions. Analysis of fMRI data from this study may also reveal that the iTBS affects the brain function rather than structure. However, even if no changes are observed by neuroimaging, patients may still experience improvements in their symptoms or affect. The analysis of patient-reported outcome measures is not in the scope of this thesis but is clearly a critical piece of advancing the iTBS intervention for chronic pain. Therefore, while the analysis presented in this thesis does not provide evidence for the mechanistic hypothesis that iTBS can structurally influence the DPMS, this result does not rule out other mechanisms of the intervention. Depending on the results from patient-reported outcomes (which are more clinically pertinent), iTBS may still be a viable and effective treatment for patients with chronic pain conditions.

3. Choice of methods and future directions

3.1. Potential alternative methods

This section discusses alternative methods that could have been applied in this analysis. Using different methods at certain steps could have changed the results obtained. The definition of ROIs is a step with significant flexibility, and an alternative implementation of the Glasser atlas should be discussed. The Glasser atlas is meant to be implemented as a classifier algorithm to be used on datasets with T1, T2, a b0 field map, and over 30 minutes of resting-state fMRI data. It uses these data as an input and outputs a segmentation based on the multimodal features of the data [Glasser 2016]. While fMRI was collected for this study, the duration of fMRI data acquired at each time point was less than 30 minutes. Therefore, a volumetric version of the atlas in MNI space was used instead. However, it may be possible to concatenate the rs-fMRI data from the pre- and post-iTBS scans together to create a dataset of sufficient length for the classifier. This would allow utilization of the Glasser classifier in its intended form and generate segmentations in native T1 space. Recall that the advantage of using the Freesurfer atlas was to circumvent the need to nonlinear register MNI-based atlases to T1 space. Using the Glasser classifier would retain this advantage and is therefore a reasonable alternative to the Freesurfer parcellations. One point to consider, though, is that changes in rs-fMRI are explicitly hypothesized in the BoostCPM protocol. The validity of concatenating two datasets which are hypothesized to be different for the purpose of segmentation is debatable. On one hand, the classifier may be more tuned to detect general features that are unlikely to change between the two sessions, such as the boundaries of established networks like the default mode network and central executive network. The hypothesized changes in connectivity strengths between specific nodes may not hold much weight in the classification algorithm. On the other hand, it can be argued that the classifier should be trained on baseline rs-fMRI data, and one might expect significant changes in rs-fMRI dynamics following an accelerated and intensive brain stimulation treatment.

Further downstream in the analysis, thresholding could have been more broadly applied to guard against errant voxels. Thresholding was performed for MD to only consider data within a physiologically meaningful range (0.0005-0.001 mm²/sec). Similarly, thresholds could have been applied for MK, ODI, ISOVF, and ICVF to only consider physiologically reasonable values.

The quantitative outputs could also have been normalized to the values obtained in the control regions for each dataset. Normalization could protect against random intersession drift due to inherent experimental variability. Since it was hypothesized that no effects should be seen in the

cuneus, any drift in median dMRI outcome in this region could perhaps be attributed to intersession drift. By normalizing to a control region, dMRI outcomes would be relative to that day's baseline. However, since pre-correction significance was observed in four statistical tests in the two control regions, this approach may be less valid.

3.2. Future analyses

Additional analyses may be applied to this dataset to explore different brain regions and relationships between these results and other data from the BoostCPM trial. At a basic level, more ROIs could be added to the analysis. For example, the amygdala and mediodorsal thalamus would be pertinent ROIs to explore given their involvement in cortical pain processing. While this thesis focused on analyzing microstructural neuroplasticity in GM regions of the DPMS, further analyses could investigate changes in the WM tracts connecting these regions. In particular, the vmPFC-dmPFC tract and connections from the PFC to the ACC are of interest for descending pain modulation. The analysis of WM could be done with tract-based spatial statistics (TBSS), a common method of performing statistical tests on quantitative dMRI data. Briefly, this method would generate a WM “skeleton” on each dataset by finding the center of each WM tract. Then, these WM skeletons are spatially registered to a common space and voxel-wise statistics can be performed on the group-level. TBSS is implemented in the BRC pipeline software used to analyze this dataset and could therefore be performed within the same framework of methods.

Changes in dMRI metrics could also be analyzed on a voxel-wise basis instead of summarizing the outcomes within ROIs. This would entail transforming the quantitative maps of the pre- and post-*ï*TBS time points to a common space and subtracting to generate a difference map. This kind of visualization could reveal whether there are systematic increases or decreases in any region. If the maps were all transformed to the standard MNI152 space, group-level statistics could be performed on the difference maps. Note that voxel-wise analysis relies on highly accurate spatial registrations.

Finally, it would be useful to correlate the quantitative dMRI outcomes with patient-reported outcomes. The baseline values and pre-post changes in each metric could be correlated with, for example, change in pain scores. It is typical that not all patients will respond to a treatment, and this certainly the case for brain stimulation therapies. Thus, correlating insights from dMRI with patient outcomes could reveal a microstructure feature that seems to predict which patients respond to treatment. It could also reveal that patient improvements are associated with certain changes in microstructure, an effect that could be obscured when considering the whole cohort. However, it should be noted that these forms of “responder analyses” can be fraught with methodological issues [Cook 2023]. If applied properly, though, exploring the relationship between neuroimaging outcomes and patient outcomes is a crucial next step with potentially significant clinical implications.

4. Conclusion

This analysis of dMRI data from patients with chronic knee pain did not reveal significant microstructural changes in GM hubs of the DPMS in response to an accelerated, connectivity-guided *ï*TBS treatment targeting the IDLPFC. Structural and dMRI data from the BoostCPM pilot clinical trial were processed with a modern in-house pipeline, and the results of each step were carefully assessed. Three dMRI models were applied to the processed data: the diffusion tensor model, diffusion kurtosis model, and NODDI model. Within-group and between-group statistical

tests were performed. While six tests reached statistical significance when performed independently, none survived p-value correction for multiple tests. Furthermore, when holistically considering observed shortcomings in data processing, the high variance in each outcome metric, and lack of consistency and mechanistic explanations for the six pre-correction significant results, it is perhaps most reasonably concluded that these are spurious results. The analysis presented does not support that iTBS induces microstructural change. However, further analyses as discussed in section 4 may yet reveal that iTBS does have an effect on neural structure. This thesis presented a purely mechanistic analysis to supplement results from patient-reported outcomes, which are of far greater relevance for the future viability of iTBS as a clinical treatment for chronic pain in knee osteoarthritis.

V. Appendices

Appendix A: Understanding the off-diagonal elements in the diffusion tensor

The necessity of the off-diagonal elements of the diffusion tensor $\mathbf{D}$ can be understood by considering the following example. Consider a voxel with a single axon passing through at an angle which is exactly 45 degrees to the x and y axes.




Figure 23. An axon oriented 45° to the x and y axes.

A diffusion-weighted image of this voxel is acquired, with gradients pulsed in the x and y directions. The resulting ADCs measured in each of these directions will be equivalent, accurately indicating that diffusion in the x and y directions is equally restricted. This might mislead one to conclude that diffusion in the voxel is perfectly isotropic when, in fact, it is perfectly anisotropic. However, the off-diagonal element D_{xy} will capture perfect correlation between the diffusivity patterns in both directions. The reality of the diffusivity profile is made even more obvious by decomposing this diffusion tensor into its eigenvectors and eigenvalues.

$$\mathbf{D} = \begin{pmatrix} D_{xx} & D_{xy} \\ D_{yx} & D_{yy} \end{pmatrix} = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}$$

The eigenvector decomposition clearly shows that all diffusion in the voxel occurs in one direction ($\lambda_1=1$) with 0 diffusion in the secondary direction ($\lambda_2=0$). The diffusion is purely anisotropic.

Appendix B: TMS protocol parameters

To expand upon the basic TMS paradigm, many parameters can be modified and must be considered when designing a TMS study. These include coil geometry, stimulation intensity, pulse shape, stimulation pattern, and total dosage [Caulfield 2022]. Improving upon the simple geometry of the original circular TMS coil, Ueno and colleagues introduced the figure-eight coil in 1988 [Ueno 1988]. This design utilizes two circular coils with current running in opposite directions. The resulting eddy currents induced in the brain by the two coils intersect and sum at their midpoint, thereby focusing the stimulation at that point. The figure-eight coil is commonly used today in the clinic and in research due to this specificity of stimulation. Other coil geometries also exist for specialized applications, but these are beyond the scope of this thesis.

The parameter of stimulation intensity (SI) indicates the magnitude of the magnetic pulse. Conceptually, this is most readily described with the peak magnetic field strength of the induced pulse (measured in T) and the rate of change of the magnetic pulse (dB/dT). In practice, however, SI is commonly described as either a percentage of maximum output or a percentage of an individual's resting motor threshold (rMT). Percentage of maximum output is relative to the hardware characteristics of the stimulator, while percent rMT scales to an individual's motor threshold. The rMT of an individual is the intensity at which stimulation of M1 induces an MEP in 50% of trials. This quantity is variable across individuals (reflecting factors such as skull thickness and cortical excitability) and is commonly used to scale stimulation intensities based on the participant. A more recent and perhaps more objective way to calculate a standardized stimulation intensity is with MRI-based electric-field modeling, which uses structural information to estimate how much energy reaches various parts of the cortex [Stenroos 2019].

The shape of a TMS pulse can be monophasic or biphasic, and its width can vary. Most TMS studies use biphasic stimulation, where current runs through the coil in both directions sequentially, inducing currents bidirectionally in the targeted area. In contrast, monophasic stimulation rectifies the pulse to a single direction. The width of the TMS pulse alters the total amount of energy delivered, with wider pulses imparting greater energy per pulse.

Stimulation pattern encompasses a large parameter space describing the timing of pulses and pauses during a TMS protocol. A key parameter within this space is stimulation frequency: the number of pulses delivered per second (expressed in Hz). Early TMS stimulators were limited by hardware to frequencies less than 1Hz. In 1990, a stimulator capable of generating patterns at >1Hz was manufactured, ushering in repetitive TMS (rTMS) protocols which are defined by this stimulation frequency cutoff. Modern stimulators can generate frequencies over 100 Hz. TMS stimulation patterns often consist of trains of pulses with pauses interspersed. For example, a 2024 study using TMS in fibromyalgia used trains of 20Hz pulses for 2 seconds separated by an inter-train interval pause of 20 seconds [Tilbor 2024].

Finally, the total dosage of a TMS protocol is the total number of pulses delivered. Some TMS interventions consist of a single session on one day, while more longitudinal designs deliver multiple TMS sessions over weeks. Recent accelerated protocols have compressed several sessions into a single day, shortening the total duration of the intervention (e.g. 10 sessions/day for 5 days instead of 1 session/day for 50 days) while delivering the same number of total pulses [Cole 2020].

The parameter space for designing a TMS study is large and altering any of the variables described here may change the effect of the intervention.

VI. Bibliography

1. Abdin, E., Chong, S. A., Vaingankar, J. A., Shafie, S., Verma, S., Luo, N., Tan, K. B., James, L., Heng, D., & Subramaniam, M. (2020). Impact of mental disorders and chronic physical conditions on quality-adjusted life years in Singapore. *Scientific Reports*, 10(1). <https://doi.org/10.1038/s41598-020-59604-0>
2. Alwardat, M., Pisani, A., Etoom, M., Carpenedo, R., Chine, E., Dauri, M., Leonardis, F., & Natoli, S. (2020). Is transcranial direct current stimulation (tDCS) effective for chronic low back pain? A systematic review and meta-analysis. *Journal of Neural Transmission*, 127, 1257–1270.
3. Ambrose, K. R., & Golightly, Y. M. (2015). Physical exercise as non-pharmacological treatment of chronic pain: Why and when. *Best Practice & Research Clinical Rheumatology*, 29(1), 120–130. <https://doi.org/https://doi.org/10.1016/j.berh.2015.04.022>
4. Andersson, J. L. R., Jenkinson, M., & Smith, S. (2007). Non-linear registration, aka Spatial normalisation FMRIB technical report TR07JA2. *FMRIB Analysis Group of the University of Oxford*, 2(1).
5. Andersson, J. L. R., & Sotiropoulos, S. N. (2015). Non-parametric representation and prediction of single-and multi-shell diffusion-weighted MRI data using Gaussian processes. *Neuroimage*, 122, 166–176.
6. Andersson, J. L. R., & Sotiropoulos, S. N. (2016). An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. *NeuroImage*, 125, 1063–1078. <https://doi.org/https://doi.org/10.1016/j.neuroimage.2015.10.019>
7. Avants, B. B., Grossman, M., & Gee, J. C. (2006). Symmetric Diffeomorphic Image Registration: Evaluating Automated Labeling of Elderly and Neurodegenerative Cortex and Frontal Lobe. In *LECT NOTES COMPUT SC* (Vol. 4057, pp. 50–57). Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/11784012_7
8. Barker, A. T., Jalinous, R., & Freeston, I. L. (1985). NON-INVASIVE MAGNETIC STIMULATION OF HUMAN MOTOR CORTEX. *The Lancet*, 325(8437), 1106–1107. [https://doi.org/https://doi.org/10.1016/S0140-6736\(85\)92413-4](https://doi.org/https://doi.org/10.1016/S0140-6736(85)92413-4)
9. Baron, R., Hans, G., & Dickenson, A. H. (2013). Peripheral input and its importance for central sensitization. *Annals of Neurology*, 74(5), 630–636. <https://doi.org/10.1002/ana.24017>
10. Basbaum, A. I., Bautista, D. M., Scherrer, G., & Julius, D. (2009). Cellular and Molecular Mechanisms of Pain. *Cell*, 139(2), 267–284. <https://doi.org/10.1016/j.cell.2009.09.028>
11. Basser, P. J., Mattiello, J., & Lebihan, D. (1994). Estimation of the Effective Self-Diffusion Tensor from the NMR Spin Echo. *Journal of Magnetic Resonance. Series B*, 103(3), 247–254. <https://doi.org/10.1006/jmrb.1994.1037>
12. Bastiani, M., Cottaar, M., Fitzgibbon, S. P., Suri, S., Alfaro-Almagro, F., Sotiropoulos, S. N., Jbabdi, S., & Andersson, J. L. R. (2019). Automated quality control for within and between studies diffusion MRI data using a non-parametric framework for movement and distortion correction. *NeuroImage (Orlando, Fla.)*, 184, 801–812. <https://doi.org/10.1016/j.neuroimage.2018.09.073>
13. Benarroch, E. E. (2012). Periaqueductal gray. *Neurology*, 78(3), 210–217. <https://doi.org/10.1212/WNL.0b013e31823fcdce>
14. Benjamini, Y., & Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1), 289–300. <http://www.jstor.org/stable/2346101>

15. Bittar, R. G., Kar-Purkayastha, I., Owen, S. L., Bear, R. E., Green, A., Wang, S., & Aziz, T. Z. (2005). Deep brain stimulation for pain relief: A meta-analysis. *Journal of Clinical Neuroscience*, 12(5), 515–519. <https://doi.org/10.1016/j.jocn.2004.10.005>
16. Boccard, S. G. J., Pereira, E. A. C., & Aziz, T. Z. (2015). Deep brain stimulation for chronic pain. *Journal of Clinical Neuroscience*, 22(10), 1537–1543. <https://doi.org/10.1016/j.jocn.2015.04.005>
17. Borsook, D., Edwards, R., Elman, I., Becerra, L., & Levine, J. (2013). Pain and analgesia: The value of salience circuits. *Progress in Neurobiology*, 104, 93–105. <https://doi.org/10.1016/j.pneurobio.2013.02.003>
18. Broderick, J. E., Keefe, F. J., Schneider, S., Junghaenel, D. U., Bruckenthal, P., Schwartz, J. E., Kaell, A. T., Caldwell, D. S., McKee, D., & Gould, E. (2016). Cognitive behavioral therapy for chronic pain is effective, but for whom? *Pain (Amsterdam)*, 157(9), 2115–2123. <https://doi.org/10.1097/j.pain.0000000000000626>
19. Brodt, S., Gais, S., Beck, J., Erb, M., Scheffler, K., & Schönauer, M. (2018). Fast track to the neocortex: A memory engram in the posterior parietal cortex. *Science (American Association for the Advancement of Science)*, 362(6418), 1045–1048. <https://doi.org/10.1126/science.aau2528>
20. Brown, R. (1828). XXVII. A brief account of microscopical observations made in the months of June, July and August 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies. *The Philosophical Magazine*, 4(21), 161–173. <https://doi.org/10.1080/14786442808674769>
21. Caulfield, K. A., & Brown, J. C. (2022). The Problem and Potential of 'TMS' Infinite Parameter Space: A Targeted Review and Road Map Forward. *Frontiers in Psychiatry*, 13. <https://www.frontiersin.org/journals/psychiatry/articles/10.3389/fpsy.2022.867091>
22. Chang, M. C., Briand, M.-M., Boudier-Revéret, M., & Yang, S. (2023). Effectiveness of transcranial alternating current stimulation for controlling chronic pain: a systematic review. *Frontiers in Neurology*, 14. <https://www.frontiersin.org/journals/neurology/articles/10.3389/fneur.2023.1323520>
23. Chen, Q., & Heinricher, M. M. (2019). Descending control mechanisms and chronic pain. *Current Rheumatology Reports*, 21, 1–7.
24. Chou, R., Turner, J. A., Devine, E. B., Hansen, R. N., Sullivan, S. D., Blazina, I., Dana, T., Bougatsos, C., & Deyo, R. A. (2015). The effectiveness and risks of long-term opioid therapy for chronic pain: A systematic review for a national institutes of health pathways to prevention workshop. In *Annals of Internal Medicine* (Vol. 162, Issue 4, pp. 276–286). American College of Physicians. <https://doi.org/10.7326/M14-2559>
25. Cole, E. J., Stimpson, K. H., Bentzley, B. S., Gulser, M., Cherian, K., Tischler, C., Nejad, R., Pankow, H., Choi, E., Aaron, H., Espil, F. M., Pannu, J., Xiao, X., Duvio, D., Solvason, H. B., Hawkins, J., Guerra, A., Jo, B., Raj, K. S., ... Williams, N. R. (2020). Stanford Accelerated Intelligent Neuromodulation Therapy for Treatment-Resistant Depression. *American Journal of Psychiatry*, 177(8), 716–726. <https://doi.org/10.1176/appi.ajp.2019.19070720>
26. Cook, C. E., Bejarano, G., Reneker, J., Vigotsky, A. D., & Riddle, D. L. (2023). Responder analyses: a methodological mess. *Journal of Orthopaedic & Sports Physical Therapy*, 53(11), 652–654.
27. Corder, G., Ahanonu, B., Grewe, B. F., Wang, D., Schnitzer, M. J., & Scherrer, G. (2019). An amygdalar neural ensemble that encodes the unpleasantness of pain. *Science (American Association for the Advancement of Science)*, 363(6424), 276–281. <https://doi.org/10.1126/science.aap8586>
28. Cottam, W. J., Iwabuchi, S. J., Drabek, M. M., Reckziegel, D., & Auer, D. P. (2018). Altered connectivity of the right anterior insula drives the pain connectome changes in chronic knee osteoarthritis. *Pain*, 159(5), 929–938.

29. Cruccu, G., Garcia-Larrea, L., Hansson, P., Keindl, M., Lefaucheur, J., Paulus, W., Taylor, R., Tronnier, V., Truini, A., & Attal, N. (2016). EAN guidelines on central neurostimulation therapy in chronic pain conditions. *European Journal of Neurology*, 23(10), 1489–1499.
30. Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., & Killiany, R. J. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage (Orlando, Fla.)*, 31(3), 968–980. <https://doi.org/10.1016/j.neuroimage.2006.01.021>
31. Destrieux, C., Fischl, B., Dale, A., & Halgren, E. (2010). Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *NeuroImage (Orlando, Fla.)*, 53(1), 1–15. <https://doi.org/10.1016/j.neuroimage.2010.06.010>
32. Dimitroulas, T., Duarte, R. v, Behura, A., Kitas, G. D., & Raphael, J. H. (2014). Neuropathic pain in osteoarthritis: A review of pathophysiological mechanisms and implications for treatment. *Seminars in Arthritis and Rheumatism*, 44(2), 145–154. <https://doi.org/https://doi.org/10.1016/j.semarthrit.2014.05.011>
33. Drabek, M., Hodkinson, D., Horvath, S., Millar, B., Pszczolkowski Parraguez, S., Tench, C. R., Tanasescu, R., Lankappa, S., Morriss, R., Walsh, D., & Auer, D. P. (2023). Brain connectivity-guided, Optimised theta burst transcranial magnetic stimulation to improve Central Pain Modulation in knee Osteoarthritis Pain (BoostCPM): protocol of a pilot randomised clinical trial in a secondary care setting in the UK. *BMJ Open*, 13(10), e073378. <https://doi.org/10.1136/bmjopen-2023-073378>
34. Einstein, A. (1906). On the theory of the Brownian movement. *Ann. Phys*, 19(4), 371–381.
35. Fayaz, A., Croft, P., Langford, R. M., Donaldson, L. J., & Jones, G. T. (2016). Prevalence of chronic pain in the UK: a systematic review and meta-analysis of population studies. *Open*, 6, 10364. <https://doi.org/10.1136/bmjopen-2015>
36. Field, R., Pourkazemi, F., Turton, J., & Rooney, K. (2021). Dietary Interventions Are Beneficial for Patients with Chronic Pain: A Systematic Review with Meta-Analysis. *Pain Medicine*, 22(3), 694–714. <https://doi.org/10.1093/pm/pnaa378>
37. Fields, H. (2004). State-dependent opioid control of pain. *Nature Reviews. Neuroscience*, 5(7), 565–575. <https://doi.org/10.1038/nrn1431>
38. Fingleton, C., Smart, K., Moloney, N., Fullen, B. M., & Doody, C. (2015). Pain sensitization in people with knee osteoarthritis: a systematic review and meta-analysis. *Osteoarthritis and Cartilage*, 23(7), 1043–1056. <https://doi.org/https://doi.org/10.1016/j.joca.2015.02.163>
39. Fischl, B. (2012). FreeSurfer. *NeuroImage*, 62(2), 774–781. <https://doi.org/https://doi.org/10.1016/j.neuroimage.2012.01.021>
40. Fonov, V. S., Evans, A. C., McKinstry, R. C., Alml, C. R., & Collins, D. L. (2009). Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage (Orlando, Fla.)*, 47, S102–S102. [https://doi.org/10.1016/S1053-8119\(09\)70884-5](https://doi.org/10.1016/S1053-8119(09)70884-5)
41. Fox, M. D., Buckner, R. L., White, M. P., Greicius, M. D., & Pascual-Leone, A. (2012). Efficacy of transcranial magnetic stimulation targets for depression is related to intrinsic functional connectivity with the subgenual cingulate. *Biological Psychiatry*, 72(7), 595–603.
42. Gaskin, D. J., & Richard, P. (2012). The economic costs of pain in the United States. *Journal of Pain*, 13(8), 715–724. <https://doi.org/10.1016/j.jpain.2012.03.009>
43. Goldberg, D. S., & McGee, S. J. (2011). Pain as a global public health priority. *BMC Public Health*, 11. <https://doi.org/10.1186/1471-2458-11-770>

44. Gorgolewski, K. J., Auer, T., Calhoun, V. D., Craddock, R. C., Das, S., Duff, E. P., Flandin, G., Ghosh, S. S., Glatard, T., & Halchenko, Y. O. (2016). The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Scientific Data*, 3(1), 1–9.
45. Greve, D. N., & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, 48(1), 63–72. <https://doi.org/https://doi.org/10.1016/j.neuroimage.2009.06.060>
46. Guillemaud, R., & Brady, M. (1997). Estimating the bias field of MR images. *IEEE Transactions on Medical Imaging*, 16(3), 238–251. <https://doi.org/10.1109/42.585758>
47. Hanson, L. G. (2008). Is quantum mechanics necessary for understanding magnetic resonance? *Concepts in Magnetic Resonance Part A*, 32A(5), 329–340. <https://doi.org/https://doi.org/10.1002/cmr.a.20123>
48. Heinricher, M. M., Morgan, M. M., Tortorici, V., & Fields, H. L. (1994). Disinhibition of off-cells and antinociception produced by an opioid action within the rostral ventromedial medulla. *Neuroscience*, 63(1), 279–288. [https://doi.org/10.1016/0306-4522\(94\)90022-1](https://doi.org/10.1016/0306-4522(94)90022-1)
49. Helm, S., Shirsat, N., Calodney, A., Abd-Elseyed, A., Kloth, D., Soin, A., Shah, S., & Trescot, A. (2021). Peripheral Nerve Stimulation for Chronic Pain: A Systematic Review of Effectiveness and Safety. *Pain and Therapy*, 10(2), 985–1002. <https://doi.org/10.1007/s40122-021-00306-4>
50. Hilton, L., Hempel, S., Ewing, B. A., Apaydin, E., Xenakis, L., Newberry, S., Colaiaco, B., Maher, A. R., Shanman, R. M., Sorbero, M. E., & Maglione, M. A. (2017). Mindfulness Meditation for Chronic Pain: Systematic Review and Meta-analysis. *Annals of Behavioral Medicine*, 51(2), 199–213. <https://doi.org/10.1007/s12160-016-9844-2>
51. Huang, J., Gadotti, V. M., Chen, L., Souza, I. A., Huang, S., Wang, D., Ramakrishnan, C., Deisseroth, K., Zhang, Z., & Zamponi, G. W. (2019). A neuronal circuit for activating descending modulation of neuropathic pain. *Nature Neuroscience*, 22(10), 1659–1668. <https://doi.org/10.1038/s41593-019-0481-5>
52. Huang, Y.-Z., Edwards, M. J., Rounis, E., Bhatia, K. P., & Rothwell, J. C. (2005). Theta Burst Stimulation of the Human Motor Cortex. *Neuron (Cambridge, Mass.)*, 45(2), 201–206. <https://doi.org/10.1016/j.neuron.2004.12.033>
53. Jensen, J. H., Helpner, J. A., Ramani, A., Lu, H., & Kaczynski, K. (2005). Diffusional kurtosis imaging: The quantification of non-gaussian water diffusion by means of magnetic resonance imaging. *Magnetic Resonance in Medicine*, 53(6), 1432–1440. <https://doi.org/10.1002/mrm.20508>
54. Julius, D., & Basbaum, A. I. (2001). Molecular mechanisms of nociception. *Nature*, 413(6852), 203–210. <https://doi.org/10.1038/35093019>
55. Keuken, M. C., Bazin, P.-L., Crown, L., Hootsmans, J., Laufer, A., Muller-Axt, C., Sier, R., van der Putten, E. J., Schaefer, A., Turner, R., & Forstmann, B. U. (2014). Quantifying inter-individual anatomical variability in the subcortex using 7 T structural MRI. *NeuroImage (Orlando, Fla.)*, 94, 40–46. <https://doi.org/10.1016/j.neuroimage.2014.03.032>
56. Kuijpers, T., van Middelkoop, M., Rubinstein, S. M., Ostelo, R., Verhagen, A., Koes, B. W., & van Tulder, M. W. (2011). A systematic review on the effectiveness of pharmacological interventions for chronic non-specific low-back pain. *European Spine Journal*, 20(1), 40–50. <https://doi.org/10.1007/s00586-010-1541-4>
57. le Bars, D., Dickenson, A. H., & Besson, J.-M. (1979). Diffuse noxious inhibitory controls (DNIC). I. Effects on dorsal horn convergent neurones in the rat. *PAIN*, 6(3). https://journals.lww.com/pain/fulltext/1979/06000/diffuse_noxious_inhibitory_controls__dnic___i_.4.aspx

58. le Bars, D., Villanueva, L., Willer, J. C., & Bouhassira, D. (1992). Diffuse Noxious Inhibitory Controls (DNIC) in animals and in man. *Acupuncture in Medicine: Journal of the British Medical Acupuncture Society*, 9(2), 47–56. <https://doi.org/10.1136/aim.9.2.47>
59. le Bihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., & Laval-Jeantet, M. (1986). MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2), 401–407.
60. Lefaucheur, J.-P. (2008). Use of repetitive transcranial magnetic stimulation in pain relief. *Expert Review of Neurotherapeutics*, 8(5), 799–808.
61. Legrain, V., Iannetti, G. D., Plaghki, L., & Mouraux, A. (2011). The pain matrix reloaded: A salience detection system for the body. *Progress in Neurobiology*, 93(1), 111–124. <https://doi.org/10.1016/j.pneurobio.2010.10.005>
62. Levy, R. M., Deer, T. R., & Henderson, J. (2010). Intracranial neurostimulation for pain control: a review. *Pain Physician*, 13(2), 157.
63. Lewis, G. N., Rice, D. A., & McNair, P. J. (2012). Conditioned Pain Modulation in Populations With Chronic Pain: A Systematic Review and Meta-Analysis. *The Journal of Pain*, 13(10), 936–944. <https://doi.org/https://doi.org/10.1016/j.jpain.2012.07.005>
64. Lindsay, N. M., Chen, C., Gilam, G., Mackey, S., & Scherrer, G. (2021). Brain circuits for pain and its treatment. *Science Translational Medicine*, 13(619), eabj7360–eabj7360. <https://doi.org/10.1126/scitranslmed.abj7360>
65. Linnman, C., Moulton, E. A., Barmettler, G., Becerra, L., & Borsook, D. (2012). Neuroimaging of the periaqueductal gray: State of the field. *NeuroImage (Orlando, Fla.)*, 60(1), 505–522. <https://doi.org/10.1016/j.neuroimage.2011.11.095>
66. Lu, C., Yang, T., Zhao, H., Zhang, M., Meng, F., Fu, H., Xie, Y., & Xu, H. (2016). Insular Cortex is Critical for the Perception, Modulation, and Chronification of Pain. *Neuroscience Bulletin*, 32(2), 191–201. <https://doi.org/10.1007/s12264-016-0016-y>
67. Lubejko, S. T., Graham, R. D., Livrizzi, G., Schaefer, R., Banghart, M. R., & Creed, M. C. (2022). The role of endogenous opioid neuropeptides in neurostimulation-driven analgesia. *Frontiers in Systems Neuroscience*, 16, 1044686. <https://doi.org/10.3389/fnsys.2022.1044686>
68. Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., Woods, R., Paus, T., Simpson, G., Pike, B., Holmes, C., Collins, L., Thompson, P., MacDonald, D., Iacoboni, M., Schormann, T., Amunts, K., Palomero-Gallagher, N., Geyer, S., ... Mazoyer, B. (2001). A Probabilistic Atlas and Reference System for the Human Brain: International Consortium for Brain Mapping (ICBM). *Philosophical Transactions: Biological Sciences*, 356(1412), 1293–1322. <http://www.jstor.org/nottingham.idm.oclc.org/stable/3067153>
69. McGrath, C. L., Kelley, M. E., Holtzheimer, P. E., Dunlop, B. W., Craighead, W. E., Franco, A. R., Craddock, R. C., & Mayberg, H. S. (2013). Toward a neuroimaging treatment selection biomarker for major depressive disorder. *JAMA Psychiatry*, 70(8), 821–829.
70. Michael, J. W.-P., Schlüter-Brust, K. U., & Eysel, P. (2010). The epidemiology, etiology, diagnosis, and treatment of osteoarthritis of the knee. *Deutsches Ärzteblatt International*, 107(9), 152.
71. Morriss, R., Briley, P. M., Webster, L., Abdelghani, M., Barber, S., Bates, P., Brookes, C., Hall, B., Ingram, L., & Kurkar, M. (2024). Connectivity-guided intermittent theta burst versus repetitive transcranial magnetic stimulation for treatment-resistant depression: a randomized controlled trial. *Nature Medicine*, 30(2), 403–413.
72. Muthulingam, J. A., Hansen, T. M., Drewes, A. M., Olesen, S. S., & Frøkjær, J. B. (2020). Disrupted functional connectivity of default mode and salience networks in chronic pancreatitis patients. *Clinical Neurophysiology*, 131(5), 1021–1029. <https://doi.org/10.1016/j.clinph.2020.01.016>

73. Nguyen, J.-P., Dixneuf, V., Esnaut, J., Moreno, A. S., Malineau, C., Nizard, J., & Lefaucheur, J.-P. (2019). The value of high-frequency repetitive transcranial magnetic stimulation of the motor cortex to treat central pain sensitization associated with knee osteoarthritis. *Frontiers in Neuroscience*, 13, 388.
74. Ning, L., Rath, Y., Barbour, T., Makris, N., & Camprodon, J. A. (2022). White matter markers and predictors for subject-specific rTMS response in major depressive disorder. *Journal of Affective Disorders*, 299, 207–214.
75. Ong, W.-Y., Stohler, C. S., & Herr, D. R. (2019). Role of the Prefrontal Cortex in Pain Processing. *Molecular Neurobiology*, 56(2), 1137–1166. <https://doi.org/10.1007/s12035-018-1130-9>
76. Ossipov, M. H., Dussor, G. O., & Porreca, F. (2010). Central modulation of pain. *The Journal of Clinical Investigation*, 120(11), 3779–3787. <https://doi.org/10.1172/JCI43766>
77. Ossipov, M. H., Morimura, K., & Porreca, F. (2014). Descending pain modulation and chronification of pain. *Current Opinion in Supportive & Palliative Care*, 8(2), 143–151. <https://doi.org/10.1097/SPC.0000000000000055>
78. Ou, Y., Akbari, H., Bilello, M., Da, X., & Davatzikos, C. (2014). Comparative Evaluation of Registration Algorithms in Different Brain Databases With Varying Difficulty: Results and Insights. *IEEE Transactions on Medical Imaging*, 33(10), 2039–2065. <https://doi.org/10.1109/TMI.2014.2330355>
79. Patetsos, E., & Horjales-Araujo, E. (2016). Treating Chronic Pain with SSRIs: What Do We Know? *Pain Research and Management*, 2016, 2020915. <https://doi.org/10.1155/2016/2020915>
80. Penfield, W. & Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain*, 60(4), 389–443.
81. Poldrack, R. A., Baker, C. I., Durnez, J., Gorgolewski, K. J., Matthews, P. M., Munafò, M. R., Nichols, T. E., Poline, J.-B., Vul, E., & Yarkoni, T. (2017). Scanning the horizon: towards transparent and reproducible neuroimaging research. *Nature Reviews Neuroscience*, 18(2), 115–126. <https://doi.org/10.1038/nrn.2016.167>
82. Price, D. D. (2000). Psychological and Neural Mechanisms of the Affective Dimension of Pain. *Science (American Association for the Advancement of Science)*, 288(5472), 1769–1772. <https://doi.org/10.1126/science.288.5472.1769>
83. Price, D. D., Greenspan, J. D., & Dubner, R. (2003). Neurons involved in the exteroceptive function of pain. *Pain*, 106(3), 215–219.
84. Reynolds, D. v. (1969). Surgery in the Rat during Electrical Analgesia Induced by Focal Brain Stimulation. *Science (American Association for the Advancement of Science)*, 164(3878), 444–445. <https://doi.org/10.1126/science.164.3878.444>
85. Rotenberg, A., Horvath, J. C., & Pascual-Leone, A. (2014). The Transcranial Magnetic Stimulation (TMS) Device and Foundational Techniques. In A. Rotenberg, J. C. Horvath, & A. Pascual-Leone (Eds.), *Transcranial Magnetic Stimulation* (pp. 3–13). Springer New York. https://doi.org/10.1007/978-1-4939-0879-0_1
86. Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., Reiss, A. L., & Greicius, M. D. (2007). Dissociable Intrinsic Connectivity Networks for Salience Processing and Executive Control. *The Journal of Neuroscience*, 27(9), 2349–2356. <https://doi.org/10.1523/JNEUROSCI.5587-06.2007>
87. Sirucek, L., Ganley, R. P., Zeilhofer, H. U., & Schweinhardt, P. (2023). Diffuse noxious inhibitory controls and conditioned pain modulation: a shared neurobiology within the descending pain inhibitory system? *Pain (Amsterdam)*, 164(3), 463–468. <https://doi.org/10.1097/j.pain.0000000000002719>

88. Somers, T. J., Keefe, F. J., Pells, J. J., Dixon, K. E., Waters, S. J., Riordan, P. A., Blumenthal, J. A., McKee, D. C., LaCaille, L., Tucker, J. M., Schmitt, D., Caldwell, D. S., Kraus, V. B., Sims, E. L., Shelby, R. A., & Rice, J. R. (2009). Pain Catastrophizing and Pain-Related Fear in Osteoarthritis Patients: Relationships to Pain and Disability. *Journal of Pain and Symptom Management*, 37(5), 863–872. <https://doi.org/https://doi.org/10.1016/j.jpainsymman.2008.05.009>
89. Soni, A., Wanigasekera, V., Mezue, M., Cooper, C., Javaid, M. K., Price, A. J., & Tracey, I. (2019). Central sensitization in knee osteoarthritis: relating presurgical brainstem neuroimaging and pain DETECT-based patient stratification to arthroplasty outcome. *Arthritis & Rheumatology*, 71(4), 550–560.
90. Stein, N., Sprenger, C., Scholz, J., Wiech, K., & Bingel, U. (2012). White matter integrity of the descending pain modulatory system is associated with interindividual differences in placebo analgesia. *PAIN*, 153(11), 2210–2217.
91. Stejskal, E. O., & Tanner, J. E. (1965). Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics*, 42(1), 288–292.
92. Stenroos, M., & Koponen, L. M. (2019). Real-time computation of the TMS-induced electric field in a realistic head model. *Neuroimage*, 203, 116159.
93. Tan, L. L., & Kuner, R. (2021). Neocortical circuits in pain and pain relief. *Nature Reviews Neuroscience*, 22(8), 458–471. <https://doi.org/10.1038/s41583-021-00468-2>
94. Tilbor, E., Hadar, A., Portnoy, V., Ganor, O., Braw, Y., Amital, H., Ablin, J., Dror, C., Bloch, Y., & Nitzan, U. (2024). TMS in combination with a pain directed intervention for the treatment of fibromyalgia - A randomized, double-blind, sham-controlled trial. *Journal of Psychiatric Research*, 170, 167–173. <https://doi.org/10.1016/j.jpsychires.2023.12.016>
95. Todd, A. J. (2010). Neuronal circuitry for pain processing in the dorsal horn. *Nature Reviews Neuroscience*, 11(12), 823–836.
96. Treede, R. D., Rief, W., Barke, A., Aziz, Q., Bennett, M. I., Benoliel, R., Cohen, M., Evers, S., Finnerup, N. B., First, M. B., Giamberardino, M. A., Kaasa, S., Korwisi, B., Kosek, E., Lavand'Homme, P., Nicholas, M., Perrot, S., Scholz, J., Schug, S., ... Wang, S. J. (2019). Chronic pain as a symptom or a disease: The IASP Classification of Chronic Pain for the International Classification of Diseases (ICD-11). In *Pain* (Vol. 160, Issue 1, pp. 19–27). Lippincott Williams and Wilkins. <https://doi.org/10.1097/j.pain.0000000000001384>
97. Tunks, E. R., Crook, J., & Weir, R. (2008). Epidemiology of Chronic Pain with Psychological Comorbidity: Prevalence, Risk, Course, and Prognosis. *The Canadian Journal of Psychiatry*, 53(4), 224–234. <https://doi.org/10.1177/070674370805300403>
98. Ueno, S., Tashiro, T., & Harada, K. (1988). Localized stimulation of neural tissues in the brain by means of a paired configuration of time-varying magnetic fields. *Journal of Applied Physics*, 64(10), 5862–5864.
99. Urien, L., & Wang, J. (2019). Top-Down Cortical Control of Acute and Chronic Pain. *Psychosomatic Medicine*, 81(9). https://journals.lww.com/psychosomaticmedicine/fulltext/2019/11000/top_down_cortical_control_of_acute_and_chronic.12.aspx
100. Veraart, J., Christiaens, D., Dai, E., Edwards, L. J., Golkov, V., Mohammadi, S., Schilling, K. G., Aarabi, M. H., Ades-Aron, B., & Adluru, N. (2022). A Data-Driven Variability Assessment of Brain Diffusion MRI Preprocessing Pipelines. *Proceedings of the Joint Annual ISMRM-ESMRMB Meeting*.
101. Verrills, P., Sinclair, C., & Barnard, A. (2016). A review of spinal cord stimulation systems for chronic pain. *Journal of Pain Research*, 9(null), 481–492. <https://doi.org/10.2147/JPR.S108884>
102. Watson, G. S. (1965). Equatorial Distributions on a Sphere. *Biometrika*, 52(1/2), 193–201. <https://doi.org/10.2307/2333824>

103. Wiech, K. (2016). Deconstructing the sensation of pain: The influence of cognitive processes on pain perception. *Science (American Association for the Advancement of Science)*, 354(6312), 584–587. <https://doi.org/10.1126/science.aaf8934>
104. Wiffen, P. J., Derry, S., Moore, R. A., Aldington, D., Cole, P., Rice, A. S. C., Lunn, M. P. T., Hamunen, K., Haanpaa, M., Kalso, E. A., & Derry, S. (2013). Antiepileptic drugs for neuropathic pain and fibromyalgia - an overview of Cochrane reviews. *Cochrane Database of Systematic Reviews*, 2019(5), CD010567–CD010567. <https://doi.org/10.1002/14651858.CD010567.pub2>
105. Willis, W. D., & Westlund, K. N. (1997). Neuroanatomy of the pain system and of the pathways that modulate pain. In *Journal of Clinical Neurophysiology* (Vol. 14, Issue 1, pp. 2–31). <https://doi.org/10.1097/00004691-199701000-00002>
106. Woolf, C. J. & Salter, M. W. (2000). Neuronal plasticity: increasing the gain in pain. *Science*, 288(5472), 1765–1768.
107. Woolf, C. J. (2011). Central sensitization: Implications for the diagnosis and treatment of pain. *Pain (Amsterdam)*, 152(3), S2–S15. <https://doi.org/10.1016/j.pain.2010.09.030>
108. Xu, H., Seminowicz, D. A., Krimmel, S. R., Zhang, M., Gao, L., & Wang, Y. (2022). Altered Structural and Functional Connectivity of Salience Network in Patients with Classic Trigeminal Neuralgia. *The Journal of Pain*, 23(8), 1389–1399. <https://doi.org/https://doi.org/10.1016/j.jpain.2022.02.012>
109. Yang, S., & Chang, M. C. (2020). Effect of Repetitive Transcranial Magnetic Stimulation on Pain Management: A Systematic Narrative Review. *Frontiers in Neurology*, 11, 114. <https://doi.org/10.3389/fneur.2020.00114>
110. Yarnitsky, D. (2010). Conditioned pain modulation (the diffuse noxious inhibitory control-like effect): its relevance for acute and chronic pain states. *Current Opinion in Anesthesiology*, 23(5). https://journals.lww.com/co-anesthesiology/fulltext/2010/10000/conditioned_pain_modulation__the_diffuse_noxious.14.aspx
111. Yu, R., Gollub, R. L., Spaeth, R., Napadow, V., Wasan, A., & Kong, J. (2014). Disrupted functional connectivity of the periaqueductal gray in chronic low back pain. *NeuroImage: Clinical*, 6, 100–108.
112. Zhang, H., Schneider, T., Wheeler-Kingshott, C. A., & Alexander, D. C. (2012). NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage (Orlando, Fla.)*, 61(4), 1000–1016. <https://doi.org/10.1016/j.neuroimage.2012.03.072>
113. Zhang, S., Zhornitsky, S., Wang, W., Le, T. M., Dhingra, I., Chen, Y., & Li, C. R. (2021). Resting state hypothalamic and dorsomedial prefrontal cortical connectivity of the periaqueductal gray in cocaine addiction. *Addiction Biology*, 26(4), e12989.
114. Zhang, Y., Brady, M., & Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*, 20(1), 45–57. <https://doi.org/10.1109/42.906424>