

Automated non-invasive brain stimulation parameter selection using Bayesian optimisation

Thomas Reed

Supervised by Professors Roi Cohen Kadosh & Mark Woolrich

A thesis presented for the degree of
Doctor of Philosophy

Michaelmas 2020

Abstract

An increasingly utilised tool in cognitive neuroscience, which is capable of drawing causal relationships between brain function and behaviour, is transcranial electrical stimulation (tES). However, this non-invasive brain stimulation technique is commonly shown to have high levels of variability when used to modulate cognitive abilities. A potential cause of this variability is the requirement for researchers to select a number of stimulation parameters such as current intensity and frequency prior to delivering stimulation to an individual. As there is a very large number of possible parameter combinations, the selection of an optimal combination with which to modulate the behaviour of interest is highly challenging. In this thesis I employ a machine learning technique called Bayesian Optimisation (BO) to automatically search through the large stimulation parameter space to identify the optimal parameters for improving cognitive performance in complex behavioural tasks. Initially, I utilise a 'standard' BO algorithm to identify optimal parameters for improving performance in a behavioural task investigating rapid instructed task learning (Chapter 3). In a subsequent chapter I validate the results of this BO algorithm in a separate study (Chapter 4). I then attempt to extend the BO algorithm to enable it to suggest personalised stimulation parameters based on an individual's personal variables. I first apply this to a sustained attention task whilst utilising a personalised variable of individual theta:beta ratio (Chapter 5). Subsequently, I apply the personalised BO (pBO) to an arithmetic task whilst considering an individual's baseline ability as the personalised variable (Chapter 6). Finally, I validate the results of the pBO algorithm in a separate study (Chapter 7). Overall, this thesis explores the use of BO algorithms to suggest parameters for neurointerventions in complex cognitive behaviours, both at a 'one-size-fits-all' and a personalised level. Whilst results from initial optimisation studies are promising, conflicting results in validation studies highlight the difficulties of applying BO to complex human behaviour and the need for validation in further work.

Acknowledgements

The past three years have been a time of academic growth culminating in this thesis. Whilst I am very proud of this work, it was made possible through the support and inspiration offered by others. First and foremost I would like to thank my supervisor Prof. Roi Cohen Kadosh. Roi, you have always been available to talk whenever I have needed help, and have been brilliant at nudging me in the right direction when I am unsure what to do. Your advice and expertise have been invaluable and I have immensely enjoyed being your student. I would also like to thank my second supervisor Prof. Mark Woolrich, who has been very welcoming and has provided me with inspiration whenever we have met.

Another truly invaluable person in the lab has been James Sheffield. Jamie, you have always been available to listen to whatever experimental problem I have had and have offered thoughtful advice and support. Your technical abilities in coding, statistics and EEG in particular are superb and you are always willing to help. Thank you. Additionally, I am deeply grateful to both Nienke van Bueren and Malin Karstens whose determination and camaraderie kept me going through long hours of testing. I would also like to thank both Jonathan Hadida and Vu Nguyen for their patience in teaching me about Gaussian processes and Bayesian optimisation and for always answering my emails with questions about them. I have been incredibly lucky to work with all of these brilliant people.

I am of course forever grateful to my parents. Without their encouragement and support none of my academic accomplishments would have been possible. A passion for learning is one of the greatest gifts I have received from them.

The last three years have also been a period of personal growth, in which I not only got married, but also had a child. Sammy, thank you for your daily kindness, compassion and support. Jessica, thank you for making the last year infinitely more enjoyable.

Contents

1 Introduction **10**
- 1.1 Transcranial electrical stimulation (tES) . . . 11
  - 1.1.1 Transcranial alternating current stimulation . . . 12
  - 1.1.2 Transcranial random noise stimulation . . . 13
- 1.2 Variability and reliability in tES . . . 16
- 1.3 Causes of variability in brain stimulation . . . 18
  - 1.3.1 Study design . . . 19
  - 1.3.2 Stimulation parameter selection . . . 22
- 1.4 Bayesian optimisation in cognitive science . . . 24
- 1.5 The purpose of this Thesis . . . 25

2 Methods **27**
- 2.1 Bayesian Optimisation . . . 27
  - 2.1.1 Burn-in phase . . . 29
  - 2.1.2 Gaussian Process Priors . . . 30
  - 2.1.3 Covariance Functions . . . 31
  - 2.1.4 Acquisition Functions . . . 33
  - 2.1.5 Importance Sampling . . . 34
  - 2.1.6 personalised Bayesian Optimisation . . . 35
- 2.2 Electroencephalography (EEG) . . . 36
  - 2.2.1 EEG pre-processing . . . 36
  - 2.2.2 EEG analysis . . . 38

3 Stimulation parameter selection using Bayesian Optimisation **40**
- 3.1 Introduction . . . 40
- 3.2 Methods . . . 43
  - 3.2.1 Behavioural Task . . . 44
  - 3.2.2 Behavioural measures . . . 45
  - 3.2.3 Stimulation and EEG . . . 46
  - 3.2.4 Bayesian Optimisation . . . 47
- 3.3 Results . . . 48

3.3.1 Simulation analysis 50
3.4 Discussion 53

4 Validating parameter selection with Bayesian Optimisation 58
4.1 Introduction 58
4.2 Methods 59
4.2.1 Participants 59
4.2.2 Stimulation 60
4.2.3 EEG processing 61
4.3 Results 61
4.3.1 Group effects of stimulation 61
4.3.2 Comparison of validation and optimisation data 64
4.3.3 Blinding Analysis 64
4.3.4 EEG analysis 66
4.4 Discussion 68

5 Personalising stimulation parameters using Bayesian optimisation for sustained attention 73
5.1 Introduction 73
5.2 Methods 77
5.2.1 Participants 77
5.2.2 Procedure 77
5.2.3 Behavioural task 78
5.2.4 Behavioural performance 79
5.2.5 Stimulation 79
5.2.6 Theta/Beta Ratio (TBR) 80
5.2.7 Personalised Bayesian optimisation (pBO) 81
5.2.8 Sensation Questionnaires 82
5.3 Results 82
5.3.1 Performance of the pBO algorithm 82
5.3.2 Sensation analysis 84
5.4 Discussion 85

6 Personalising stimulation parameters using Bayesian optimisation for arithmetic skills 89
6.1 Introduction 89
6.2 Methods 91
6.2.1 Behavioural Task 91
6.2.2 Measurement of Behavioural Performance 92
6.2.3 Experimental Procedure 93
6.2.4 pBO Algorithm 94

6.2.5 Transcranial Alternating Current Stimulation 94
6.2.6 Resting State-EEG Recordings and Pre-processing 95
6.2.7 Sensation Questionnaire 96
6.3 Results 97
6.3.1 Performance of the pBO algorithm 97
6.3.2 Behavioural results 98
6.3.3 Simulation analysis of pBO performance 100
6.3.4 Sensation Analysis 103
6.3.5 Baseline EEG analysis 103
6.4 Discussion 104

7 Validating personalised Bayesian Optimisation for stimulation parameter selection 107
7.1 Introduction 107
7.2 Methods 107
7.2.1 Stimulation 108
7.2.2 EEG preprocessing and analysis 109
7.3 Results 109
7.3.1 Group Results 109
7.3.2 pBO analysis 111
7.3.3 Control Analysis 112
7.3.4 EEG Analysis 114
7.4 Discussion 114

8 Discussion 117
8.1 Summary 117
8.2 Advances and Limitations 120
8.2.1 Data-driven tES parameter selection 120
8.2.2 Personalised tES parameters 121
8.2.3 Validation of BO results 123
8.3 Future work 123
8.3.1 Study design 123
8.3.2 Adapting to variability in behavioural responses 124
8.3.3 Hierarchical BO 124
8.3.4 Statistical improvements 125
8.4 Concluding remarks 125

A
A.1 Screening and consent form . . . . .
A.2 Sensation questionnaire . . . . .

List of Figures

1.1 Stimulation electrode montage 11
1.2 Stimulation profile of tACS and tRNS 13
1.3 Variability in stimulation 17

2.1 Illustration of the BO procedure 29

3.1 A summary of the RITL task 45
3.2 Performance predicted by the Gaussian process in the RITL task after each participant 49
3.3 Performance predicted by the GP at each stimulation parameter combination 50
3.4 Simulation examining the effect size needed to reliably find the optimal stimulation parameter with Bayesian optimisation and random search algorithms 52

4.1 Results of the one-tailed sequential Bayesian t-tests showing no effect of tRNS on transfer 62
4.2 Drift rates during both days of the experiment 63
4.3 Performance predicted by the GP with the validation data added . 65
4.4 parietal theta power at different transfer distances 67

5.1 pBO overview 75
5.2 Sustained attention paradigm 78
5.3 pBO attentions surface 83
5.4 pBO attentions scatter plots 84

6.1 Arithmetic paradigm 92
6.2 Best predicted performance of the pBO model at each iteration . . 97
6.3 pBO arithmetic surfaces 98
6.4 pBO best frequency and current 99
6.5 Maximum predicted performance at different baseline abilities . . . 100
6.6 Comparison of pBO, BO and random sampling 102

7.1 pBO validation group results 110
7.2 pBO prediction vs validation comparison 111
7.3 Interactions between perceived and actual stimulation condition affecting performance 113

List of Tables

5.1 Sensation analysis in the attention pBO study 85

6.1 Sensation analysis in the arithmetic pBO study 103

Nomenclature

Stimulation

tES Transcranial Electrical Stimulation

$tDCS$ Transcranial Direct Current Stimulation

$tACS$ Transcranial Alternating Current Stimulation

$tRNS$ Transcranial Random Noise Stimulation

TMS Transcranial Magnetic Stimulation

Bayesian Optimisation

BO Bayesian Optimisation

pBO Personalised Bayesian Optimisation

GP Gaussian Process

PI Potential Improvement acquisition function

EI Expected Improvement acquisition function

UCB Upper Confidence Bound acquisition function

Neuroimaging

EEG Electroencephalography

$fMRI$ Functional Magnetic Resonance Imaging

MEG Magnetoencephalography

$DLPFC$ Dorsolateral Prefrontal Cortex

PPC Posterior Parietal Cortex

PSD Power Spectral Density

FFT Fast Fourier Transform

$rs - EEG$ resting state EEG

Chapter 1

Introduction

Throughout history, knowledge of the human brain and the understanding of its functioning, has been enabled by the development of new methodological techniques. From early examinations of individuals with neurological conditions and brain injuries to modern day methods, neuroscience has developed a suite of tools to link brain function to behaviour. The majority of these modern techniques such as magnetic resonance imaging (MRI), electroencephalography (EEG) and magnetoencephalography (MEG) rely on drawing correlations between behaviour and the brain. However, a subset of these modern techniques are able to draw causal connections between cognitive processes and behaviour by enabling the modulation of brain function and observation of the resulting behavioural changes. These techniques can include invasive forms of stimulation as well as non-invasive forms such as transcranial magnetic stimulation (TMS), transcranial focused ultrasound stimulation (tFUS) and transcranial electrical stimulation (tES). In this thesis I focus on tES to modulate brain function and therefore human behaviour.

1.1 Transcranial electrical stimulation (tES)

Transcranial electrical stimulation (tES) is a non-invasive form of brain stimulation reliant on the administration of an electrical current to the scalp via electrodes (Figure 1.1). The applied current is capable of penetrating the skull and influencing the underlying neuronal populations. The mechanism by which this is achieved varies between different types of stimulation, of which there are three main types; transcranial direct current stimulation (tDCS), transcranial alternating current stimulation (tACS) and transcranial random noise stimulation (tRNS). Only tACS and tRNS are utilised in this thesis.

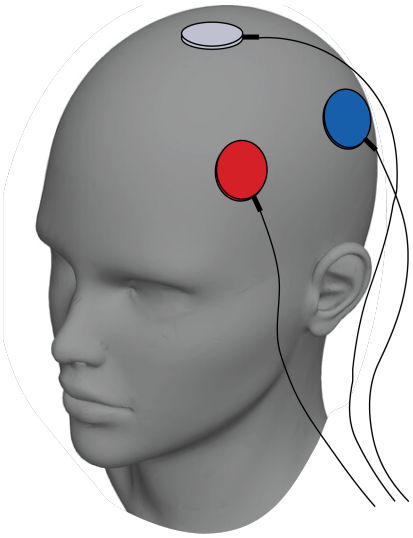


Figure 1.1: An example of the arrangement of electrodes in tES with the stimulating electrodes over the frontal and parietal areas (red and blue electrodes) and the return over the vertex (grey electrode).

1.1.1 Transcranial alternating current stimulation

tACS utilises an alternating electrical current to entrain the firing of the neurons underlying the stimulating electrodes to an exogenous frequency (Battleday et al. 2014). The profile of the stimulating current usually takes the form of a sinusoidal wave, alternating at the desired frequency (Figure 1.2A). This oscillating frequency entrains the neurons underlying the stimulating electrodes. This was originally shown by Fröhlich and McCormick (2010) who stimulated ferrets intracranially and recorded local field potentials and multiunit activity. In doing so, they showed that not only were the neural spikes in the multiunits synchronised to the oscillatory pattern of the local field potential, but when stimulation with an alternating current they were able to elicit spiking activity which entrained to the applied frequency. Whilst a potential underlying mechanism of tACS had been demonstrated, the ability of such a weak alternating current to pass through the scalp and skull and still be able to entrain the underlying neurons was not certain. Later it was demonstrated, first in rats (Ozen et al. 2010), and then in human cadavers (Vöröslakos et al. 2018), that a proportion of the current does reach the cortex. The most recent study involving cadavers (Vöröslakos et al. 2018) demonstrated that approximately 75% of the applied tES current is lost though the shunting of the current through the scalp. However it should be noted that upon death significant biophysical changes occur within the brain tissue, limiting the applicability of these results to living tissue. Whether the current which reaches the brain is enough to induce entrainment is very hard to assess. One study reported that in epilepsy patients undergoing invasive monitoring, tACS was not able to entrain neurons in the gamma range during sleep (Lafon et al. 2017), however reports of changes in corti-

cal excitability in response to tACS suggest that it is capable of entraining neurons (Guerra, Pogosyan, et al. 2016; Feurra et al. 2013). Additionally, a recent study in primates showed that tACS does affect spike timing at intensities feasible to humans (Johnson et al. 2020). The behavioural effects of tACS have been shown to be highly frequency dependent (Kanai et al. 2008), enabling researchers to causally link neuronal oscillatory frequencies and brain regions to human behaviours (for a full description of tACS see Tavakoli and Yun 2017; Herrmann et al. 2013; Liu et al. 2018).

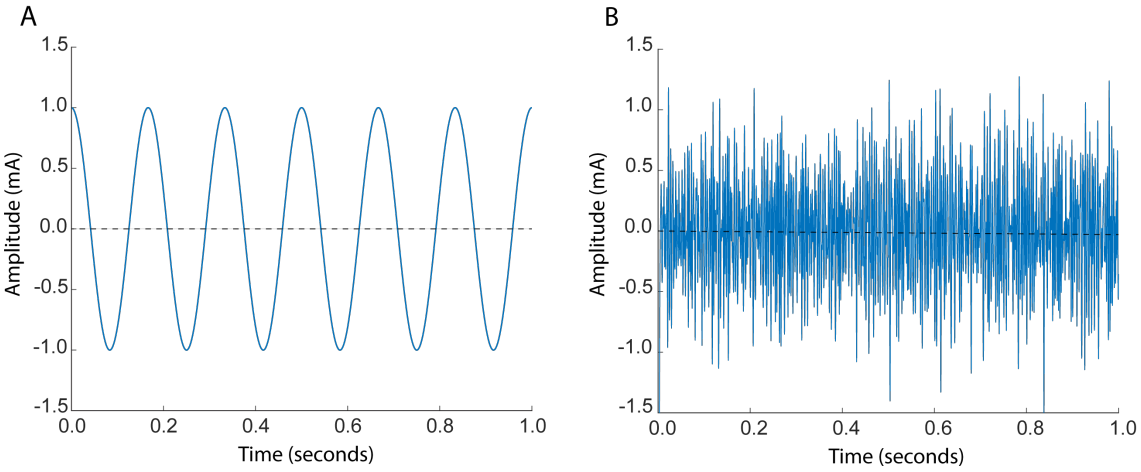


Figure 1.2: The stimulation profiles of **A)** tACS and **B)** tRNS.

1.1.2 Transcranial random noise stimulation

tRNS is a relatively recently developed tES method, being originally used in human participants in 2008 (Terney et al. 2008). tRNS is similar to tACS in that it uses an alternating current; however, instead of stimulating at a fixed frequency throughout the stimulation period, tRNS alternates at a random frequency and amplitude within a predefined range (Figure 1.2B). Stimulation frequency is nor-

mally distributed between 0.1 and 640 Hz (Terney et al. 2008), although it can be divided into either low- (0.1–100 Hz) or high- (101-640 Hz) frequency stimulation (Fertonani, Pirulli, and Miniussi 2011). The mechanistic underpinnings of tRNS are still unclear, however there are several theories of how it could work:

- Stochastic resonance - In a system in which an oscillating sub-threshold signal exists, the addition of noise may be beneficial, as it pushes the signal over the threshold. Importantly, stochastic resonance predicts an inverted-U-shape in response to varying the level of noise, as too little noise would not push the signal over the threshold, whilst too much noise would saturate the signal and the threshold would be constantly surpassed. This is clearly illustrated in the visual system in which adding noise to visual stimuli has been shown to modify perception in an inverted-U-shape (van der Groen and Wenderoth 2016). In a similar manner, adding noise directly to the occipital cortex during perceptual tasks, using tRNS at different intensities, has been shown to deliver similar results (van der Groen and Wenderoth 2016), leading researchers to propose a stochastic resonance mechanism for the action of tRNS. Whilst stochastic resonance is most easily demonstrated in the occipital cortex, it is proposed to act throughout the brain.

- Ripple oscillations - tRNS could modulate high frequency 'ripples' between 140–220Hz. These ripples originate in the hippocampus and enable the transfer of information to the neocortex and have been shown to play a role in learning and memory (Buzsáki and Silva 2012). Indeed, the suppression of hippocampal ripples has been shown to impair learning (Ego-Stengel and Wilson 2010) and memory (Girardeau et al. 2009) in the rat. In humans,

the use of 140Hz tACS over M1 has been found to increase corticospinal excitability to the hand, whilst 80Hz did not, and 250Hz did, but to a lesser degree (Moliadze, Antal, and Paulus 2010). tRNS covering these frequencies (100–640Hz), elicited changes in corticospinal excitability at M1 comparable to that of tACS at 140Hz (Terney et al. 2008), suggesting a modulatory effect of tRNS on ripple oscillations.

- Repeated subthreshold stimulation - repeated subthreshold stimulation via tRNS might prevent homeostasis of the brain's systems, or interact with the activity of the involved neurons, potentially enhancing performance (Fertonani, Pirulli, and Miniussi 2011).

These three proposed mechanisms of tRNS each have their own merits. Whilst it is possible that only one of them drives the effects seen from tRNS, I believe it is most likely that a combination of all/some of them is responsible for the observed effects of tRNS. Specifically, I think it is likely that repeated subthreshold stimulation causes changes in cortical activity whenever tRNS is applied, however whether these cortical changes are enough to contribute to changes in behaviour will depend on numerous factors including the location, intensity and frequency of stimulation and the behaviour being investigated. The involvement of ripple oscillations in tRNS effects is also plausible, however this is likely confined to behaviours which involve the hippocampus more heavily such as learning and memory. Additionally, stochastic resonance clearly has an effect in occipital cortex function, and likely has an effect elsewhere, however it is particularly hard to disentangle the expected signs of these mechanisms in more complex behaviours. This difficulty makes it extremely difficult to validate the mechanistic underpinnings of tRNS without first

conducting further in vitro work to discern the neuronal mechanisms, and then further whole-brain studies to validate this.

1.2 Variability and reliability in tES

Despite being an established research tool, an ever growing number of publications utilising tES have revealed a large amount of variability in the efficacy of stimulation, with the application of identical stimulation protocols yielding very large inter-and intra-individual variability (B. Krause and Cohen Kadosh 2014; Chew, Ho, and Loo 2015). This has resulted in a lack of reliability in experimental results, for example, Wiethoff, Hamada, and Rothwell (2014) elicited motor-evoked potentials (MEPs) after participants had undergone 10 minutes of tDCS delivered at 2mA. They found that in approximately 50% of participants, tDCS had a facilitatory effect on MEP amplitudes, whereas there was minimal or no effect of tDCS in the remaining participants. This variability was clearly summarised by a recent meta-analysis of the effects of tDCS (Guerra, López-Alonso, et al. 2020) which found a large number of studies reporting significant intra- and inter-individual variation, with the proportion of participants categorised as 'responders' to stimulation ranging from 20-90% (Figure 1.3). Whilst this meta–analysis investigates variability in tDCS, it is still illustrative of variability in tES as a whole.

This lack of reliability leads to difficulty in replicating results. For example, Veniero et al. (2017) delivered parietal alpha tACS to 20 participants whilst they completed a line bisection task. The results showed the expected changes in spatial attention bias, however upon repeating the experiment with another independent group of 20 participants, these effects were not found. A similar lack of replication

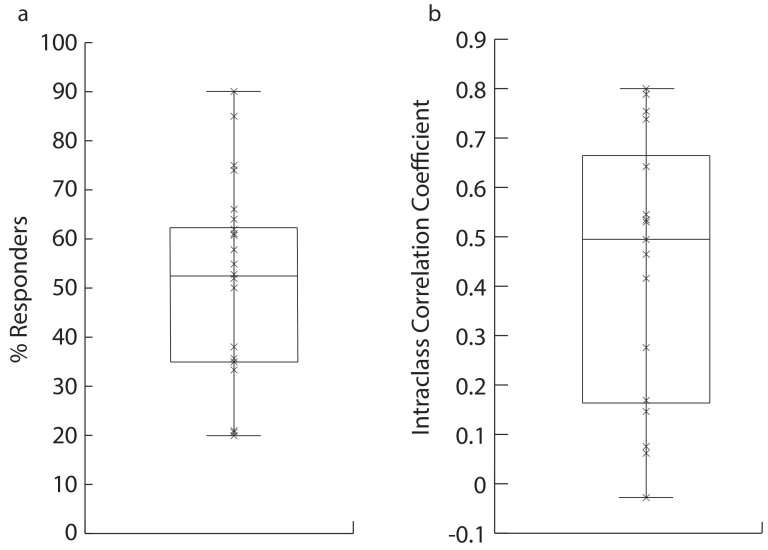


Figure 1.3: **Results of inter-and intra-individual variation in tES. a)** Percentage responders to tES in a wide range of behavioural paradigms, demonstrating inter-individual variability. **b)** Intraclass correlation coefficients in response to tES techniques in a variety of behavioural paradigms, demonstrating intra-individual variability. Data from Guerra et al.(2020).

was found in another study, which conducted 4 experiments that used tACS to modulate alpha power (Clayton, Yeung, and Cohen Kadosh 2019). Similar studies (Costa et al. 2015; Horvath, Carter, and Forte 2016; Koenigs et al. 2009; Learmonth et al. 2017; Vannorsdall et al. 2016) have also shown a lack of behavioural replicability when using tES.

When combined with failed replications, the large degree of variability observed in tES results has elicited debate over the reliability of tES as a research and clinical tool, with some researchers questioning the effects of tES on the brain and behaviour (Antal, Keeser, et al. 2015; Horvath 2015; Polanía, Nitsche, and Ruff 2018; Price and Hamilton 2015). These concerns highlight not only the importance of sufficient statistical power when using tES, but also the need to optimise tES protocols to consider individual variability and provide robust results. Currently however, the

best ways in which tES protocols can be refined is unclear. In the following sections, I describe and discuss the methods that are currently being utilised by researchers, as well as point to future ways in which reliability can be improved.

1.3 Causes of variability in brain stimulation

There are a wide range of person-specific variables that can cause variability in tES results (for review see Ridding and Ziemann 2010), including gender (Chaieb, Antal, and Paulus 2008; Fumagalli et al. 2010; Russell et al. 2017), age (Bartrés-Faz, Rossi, and Pascual-Leone 2016) and genetics (Antal, Chaieb, et al. 2014; Teo et al. 2014). Additionally, stimulation parameters such as electrode location and orientation (Wörsching et al. 2018), operator training and insufficient protocol reporting (Bikson, Brunoni, et al. 2018) clearly affect the results of tES studies. Crucially, stimulation effects have been shown to be current, frequency and phase dependent (Esmaeilpour et al. 2018; Nakazono et al. 2016), as well being altered by the time between stimulation and a target behaviour (Molero-Chamizo et al. 2018).

These factors not only affect how well an individual responds to tES, but also the 'dosage' of tES that is best for an individual, i.e. the "electrode montage and waveform (including shape, pulse width, polarity, frequency, duration, and amplitude of current as well as number and frequency of sessions)" (Bikson, Esmaeilpour, et al. 2019). Person-specific variables can play a significant role in determining the optimal dosage for an individual. For example, physiological differences such as the distance between the surface of the scalp and the underlying grey matter (Davis 2020) will affect diffusion of the tES current, and anatomical differences between individuals will reduce the accuracy of targeting key structures in the cortex (Radecke

et al. 2019). Other inter–individual differences which can influence tES dosage include: age, through reductions in cortical excitability with age (Bartrés-Faz, Rossi, and Pascual-Leone 2016), gender through differences in hormonal changes (Inghilleri et al. 2004) and handedness through hemispheric specificity. In order to combat the risk of variance that these factors pose, researchers will often either statistically control for these factors in their analysis or limit the sample of participants they use, for example to right-handed participants or within a specific age range. Intra–individual differences also influence tES dosage. For example, it has been clearly shown that the effects of tES, and therefore the dosage required, are brain-state dependent (Cohen Kadosh, Soskic, et al. 2010; Feurra et al. 2013; Neuling, Rach, and Herrmann 2013). Brain-state can be influenced through the use of behavioural tasks, however as intra–individual variables by their very nature change with time, the challenge of continuously monitoring them and adjusting dosage accordingly is a significant one. Clearly the modulation of tES dosage by an individual's variables is a complex area which requires further study. Therefore, for simplicity, this thesis is limited to investigating the modulation of tES current intensity and frequency. The following sections expand on two key causes of variability that are explored in this thesis: study design and stimulation parameter selection.

1.3.1 Study design

A key component affecting the validity of any set of experimental results is the study design. This is highlighted in psychological and neuroscientific research by the recent movement to improve reproducibility in research. One commonly recurring issue is the use of an insufficiently sized sample. This is particularly prevalent

in studies utilising tES as the use of different stimulation parameters (e.g. sham and active stimulation) often requires separate experimental sessions with time between the sessions in order to exclude carry-over effects from previous stimulation sessions. This necessity adds expense to the research process and dissuades researchers from testing a larger number of participants. As small sample sizes, and therefore lower statistical power, have been clearly associated with reduced reliability in neuroscientific research (Button et al. 2013), studies lacking the appropriate sample size required to detect a true effect of stimulation or to reduce the chance of false-positive effects risk reducing the chances of producing reliable and replicable results and therefore adversely affecting the reliability of tES.

In order to account for the lack of reliability between participants in tES research (Chew, Ho, and Loo 2015; B. Krause and Cohen Kadosh 2014; Wiethoff, Hamada, and Rothwell 2014), some researchers have taken the approach of dividing participants into 'responders' and 'non-responders' during their analysis (for review see Pellegrini, Zoghi, and Jaberzadeh 2018). One study found that after 13 minutes of anodal tDCS over the motor cortex in two separate sessions, only 60% and 64% of subjects exhibited the expected increase in cortical excitability (López-Alonso et al. 2015). In a very similar experiment, Strube et al. (2016) used cluster analysis of participants' changes in cortical excitability to both anodal and cathodal tDCS to demonstrate a bidirectional response in which 61% of participants responded to anodal stimulation and 53% responded to cathodal stimulation. Interestingly, regardless of stimulation condition, participants who were clustered for increases in excitability in response to stimulation, showed higher measures of baseline global cortico-spinal excitability. This result led the authors to suggest a tentative link between baseline excitability and the likelihood of developing increases in excitability

following tDCS. Although more research is needed to be confident of this finding (for a recent example with tRNS see Sheffield et al. 2020), links between baseline measures and changes caused by tES methods are invaluable in attempting to predict which individuals will respond well to stimulation. Whilst the studies previously mentioned utilised measures of cortical excitability to categorise responders and non-responders, these observed differences can manifest at the behavioural level (Looi et al. 2016; Santarnecchi et al. 2016; Sarkar and Cohen Kadosh 2016; Tseng et al. 2012), allowing for an alternative method of classification if neurophysiological data is not available. Recently however, issues have been raised concerning the validity of the post-hoc division of subjects into groups, with van de Ruit and Grey (2019) demonstrating that large type I error is associated with this practice.

When designing tES studies a key concept researchers must be aware of and account for is the occurrence of carry-over effects. Carry-over effects are when the neural or behavioural changes induced by one block or session of tES carry over into the next. This can occur when the tES blocks are carried out too close together. Unfortunately the degree to which tES effects last is under-explored and currently unclear. Whilst it has been shown that the polarising effects of tDCS on the resting membrane potential last approximately 90 minutes following 13 minutes of stimulation (Nitsche and Paulus 2001), in tACS and tRNS there is varying evidence as to the length of any carry-over effects. At very short durations of tACS, it has been shown that there are no lasting neuronal effects (Strüber et al. 2015), however delivery of tACS over a longer duration (Clancy et al. 2018) or over repeated sessions has been shown to have more lasting effects(Fresnoza et al. 2018). Similarly, both single session (Terney et al. 2008) and repeated session (Snowball et al. 2013) paradigms of tRNS have been shown to have continued effects after

the cessation of stimulation. The presence of potential carry-over effects poses a dilemma for researchers. In order to get the clearest results possible, time must be left between sessions of tES to avoid these carry-over effects. However, this is not always practical in all study designs. In such cases where this can not be done, every attempt must be made to counteract any order effects induced by the presence of carry-over effects, such as counter-balancing the order of different stimulation types and controlling for order in statistical analysis.

1.3.2 Stimulation parameter selection

As previously noted, the results of stimulation have been found to be frequency and current dependent (Esmaeilpour et al. 2018; Nakazono et al. 2016). Despite this being a known attribute of the effects of stimulation, the selection of the parameters used in stimulation studies in the literature has so far been somewhat arbitrary. In the case of current intensity, when designing a study, researchers will often use 1, 1.5 or 2mA based on at most a small number of previous studies that have had success with these parameters or based only on the tolerability of participants to the side effects of stimulation at that current intensity. In the defence of researchers, there are no clear guidelines on the appropriate current to use and the few studies that investigate the differential effects of various current intensities use only a small number of different values (Batsikadze et al. 2013), which cannot fully explain the effects of different currents. Additional complications are added to the process of selecting an appropriate current by the fact that the effects of a specific current intensity would be expected to vary depending on factors such as location and duration of stimulation, cognitive domain being targeted and the type of tES

being used, therefore limiting the transfer of knowledge between studies in different cognitive domains or at different neural locations. This problem is further complicated when researchers must select a specific frequency or phase (in the case of tACS) of stimulation. When utilising tACS and tRNS, researchers often have a specific frequency or frequency band at which they wish to stimulate in order to draw causal links between neural oscillations and behaviour. Whilst this is a perfectly reasonable approach to take in frequency selection, it is usually unknown whether the selected frequency will provide the greatest behavioural/neuronal change.

Additionally, the use of a 'one-size-fits-all' approach to parameter selection is incongruous with what we know of the brain and how it functions. Numerous studies report large variability between participants in terms of the location of specific peaks such as alpha and theta along the frequency range (Benwell et al. 2019; Rodriguez-Larios and Alaerts 2019; Moran et al. 2010). Despite this, in the case of tACS, researchers often continue to apply stimulation at one specific frequency and either fail to take into account underlying individual differences or simply include them as covariates in post-hoc analysis. Some studies attempt to personalise the stimulation by stimulating at an individual's peak frequency (e.g. peak alpha frequency), however again it is unknown whether this is truly the 'ideal' stimulation frequency for an individual.

An issue which researchers inevitably face when trying to solve this problem is the large number of combinations of stimulation parameters. For instance, if a researcher wished to stimulate using tACS they might have a range of currents from 0.1-2mA and a range of frequencies from 1-50Hz. Assuming a resolution of 0.1mA and 1Hz, that gives 1,000 possible stimulation parameter combinations to choose from. To determine the ideal stimulation parameters through testing each

combination individually would clearly be an enormous (and unrealistic) amount of work.

One recent possible advance in the selection of stimulation parameters comes in the form of a machine learning algorithm called Bayesian optimisation (Lorenz et al., 2016). Bayesian optimisation (BO) is a method in which a parameter space, in the case of stimulation the current x frequency space, may be efficiently searched to find the maximal outcome, such as a behavioural or neuronal response. Whilst a fixed search of each possible point would require a large number of samples, BO has the advantage of limiting the number of samples required whilst locating the maximal response (for a general review see Snoek, Larochelle, and Adams 2012).

1.4 Bayesian optimisation in cognitive science

BO was recently utilised in a study that attempted to select the optimal frequency and phase of tACS to induce a specific brain state in 8 participants in a closed-loop fMRI paradigm (Lorenz, Monti, et al. 2016). Whilst the optimisation algorithm was not successful on converging on the optimal stimulation parameters in this study, Lorenz and colleagues' use of BO to determine stimulation protocols offers a much more data-driven approach to stimulation parameter selection, than the commonly used practice of reverting to the protocols previously used in the literature. Further work from the same lab used BO to determine that differences in the relative phase between tACS electrodes affected phosphene perception (Lorenz, Simmons, et al. 2019). This study was run as a within-subject protocol, in that each of the 20 iterations of the BO algorithm was run within one participant, and a new BO algorithm was used for each of the participants. In this manner, stimulation

parameters were tailored to the individual in order to optimally induce phosphenes. Whilst this study is an excellent proof of principle for using BO in human research, it had a small sample size (10 participants in each of the two experiments) and was conducted using a physiological response which is more reliable than many behavioural responses. Additionally, the stimulation parameter space for phosphene induction is unlikely to have multiple optima, as is likely to be the case in behavioural responses. Additional work utilising BO has sought to identify optimal parameters to elicit MEPs using TMS, demonstrating its effectiveness (Tervo et al. 2020). However, this study utilised a very small number of participants (n=5) and found an optimum in a very simple search space, where other algorithms such as gradient descent would possibly be more suited due to the lack of multiple optima.

1.5 The purpose of this Thesis

As previously shown in this chapter, with the increase in use of neurointerventions such as tES, efficient and accurate selection of parameters is becoming increasingly vital. Crucially, researchers and clinicians must have a data-driven way of selecting these parameters. In order to meet these needs, I will build on limited existing work surrounding BO in human cognition, with the following aims:

- To utilise BO for parameter selection of tES in order to identify an optimal set of stimulation parameters for improving performance on a cognitive task.
- To attempt to reduce the variation in response to tES by personalising the stimulation parameters to the individual.
- To thoroughly validate all results of the BO algorithm to ensure reliability of

the selected tES parameters.

Chapter 2

Methods

2.1 Bayesian Optimisation

Bayesian optimisation (BO) is a technique which has previously been used in settings such as machine learning in which there exists an unknown objective function (e.g. behavioural performance) that must be maximised. Using sequential active sampling, BO aims to maximise this unknown objective function by evaluating the function with different parameters as few times as possible. Mathematically, BO can be expressed in the form of the following equation:

$$x^* = \text{argmax} f(x) \tag{2.1}$$

In which $f(x)$ is the unknown blackbox function being evaluated to find the maximum value (x^*).

In simplistic terms, BO works by training a Gaussian process (GP) model on an initial data set. The collection of this initial data is sometimes known as the burn-in phase and consists of data collected with random parameter inputs. The GP is able to estimate the output of the unknown objective function at all combi-

nations of the input parameters, forming a surrogate model. The main advantage of utilising a GP for this, is that it also provides an estimate of the uncertainty surrounding each estimate of the objective function. Therefore, it is possible to utilise an acquisition function to select which parameter combinations should be evaluated next. The acquisition function does this by balancing the exploration of parameter combinations which have not been tested (i.e. there will be a high level of uncertainty about them), against exploiting parameter combinations which demonstrate a high value of the objective function. Once the acquisition function has identified a parameter combination, it can be tested and the new data can be added and the GP can be retrained on all collected data, therefore updating the surrogate model. This process iterates until either the potential to further maximise the objective function is negligible and the algorithm converges on a maximum, or until the pre-defined stopping criteria are reached such as a maximum number of iterations (Figure 2.1).

Using this active sampling technique BO requires a smaller number of evaluations of the objective function compared to rigorously testing each parameter combination individually or using a random search algorithm. This quality makes BO a particularly useful method in situations in which each evaluation of the objective function is costly, for example the training of a neural network, and also means that BO lends itself to use in applications outside of machine learning, for instance when selecting which stimulation parameters give the largest increase in behavioural performance. In this situation, the costs of rigorously testing all possible parameter combinations include increased testing time and monetary expense, as well as participant safety and ethical concerns due to needless exposure to tES. Through the use of an active sampling technique such as BO, these costs can be

reduced.

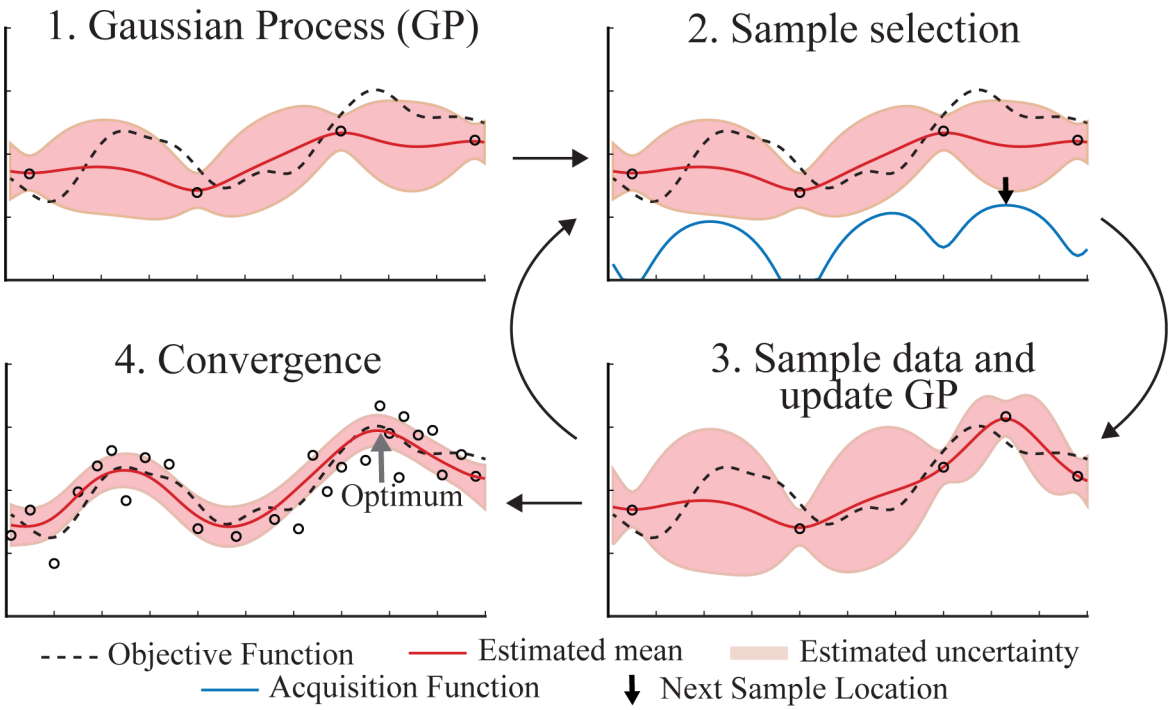


Figure 2.1: **Illustration of the BO procedure.** BO aims to find the parameter combinations (x-axis) that will maximise performance (y-axis) on an unknown objective function. This is done in 4 steps; **1)** A Gaussian process is fitted to the data, estimating the Objective function and the associated uncertainty, **2)** An acquisition function identifies the potential improvement at each parameter combination and selects the next point to sample at, **3)** the new data point is sampled, and a new Gaussian Process is fitted. **4)** This cycle continues until the potential improvement is zero at which point the algorithm is said to have converged on the optimal parameters.

2.1.1 Burn-in phase

Whilst not part of the BO algorithm, the burn-in phase is an important phase in which initial data is gathered, allowing the BO algorithm to be trained on some data before selecting new data points to test. Typically, the data points are gathered either using random input parameters, or parameters are arranged in a grid pattern

in the search space, ensuring equal distribution of initial data. In this work, a grid pattern was used in chapter 3 to measure variability in responses more thoroughly, and in chapters 5 and 6 random sampling was used due to the higher number of dimensions present in the search area. Typically, in the machine learning literature, the number of data points captured in this phase is a relatively small proportion of the overall data. However in the work presented in this thesis, substantial burn-in phases are utilised as the level of noise in the behavioural responses being measured in response to stimulation is unknown, and therefore the burn-in phase is also used to assess this and design the BO algorithm accordingly.

2.1.2 Gaussian Process Priors

Just as a Gaussian distribution is a distribution over a random variable, specified by its mean and covariance, a GP is a distribution over functions, specified by its mean function and covariance function (Brochu, Cora, and Freitas 2010). In BO, the GP models all possible functions that could be used to describe the data collected so far. By their very nature, a GP assumes that the underlying distribution of functions is Gaussian. This assumption is relevant to the applications in this thesis as it is assumed that whilst not all possible sources of noise are known, the underlying inter- and intra-individual noise in behavioural responses is normally distributed. From the infinite number of possible functions that could describe the existing data, the GP takes a mean value which serves as the model's estimate of the objective function at each parameter combination. The estimate is also known as the surrogate mean. Importantly, the GP is able to construct an estimate of the uncertainty at each point of the surrogate mean in the form of the surrogate

uncertainty.

As an example, suppose we have an individual's behavioural performance in an unspecified task f whilst they receive tES at a number of different frequencies $x_{1:k}$, in the form of a vector $f(x_{1:k})$. The GP constructs a mean vector by evaluating a mean function μ at each point of x and a covariance matrix by evaluating the covariance function $\sum$ at each pair of points within x:

$$f(x_{1:k}) = GP(\mu(x_{1:k}), \sum(x_{1:k}, x_{1:k})) \tag{2.2}$$

Using a GP it is possible to calculate the mean and variance of a normal distribution of functions at all possible values of f at x. Collectively, the estimated mean and variance measures are known as the surrogate model, with the mean corresponding to the prediction of the model at any value of x and variance corresponding to the certainty in that prediction.

2.1.3 Covariance Functions

A number of different covariance functions (also called kernels) are available to use in the GP, each encoding all of the assumptions about the form of the functions being used to model the data. For example, the degree of smoothness in the functions can be encoded by assigning a large positive correlation to points that are closer together in the input space, therefore demonstrating the prior belief that they should have more similar function values than points that are far apart (Frazier, 2018). Two key covariance functions are described below:

- Squared exponential kernel – This is a widely used kernel in the machine learning field. However it possesses very high smoothness assumptions, mak-

ing it potentially inappropriate for modelling many processes (Stein 1999 in Rasmussen & Williams, 2006).

- Matérn class of covariance functions - These functions are less smooth than the squared exponential kernel, so are therefore more applicable to noisier environments (Rasmussen and Williams 2006).

After choosing the covariance function, the hyperparameters that govern its behaviour must also be chosen. These include:

- Length scales l – This describes how smooth a function is. Small values indicate that functions can change quickly, whilst larger ones characterise functions that change slowly. This also determines how far we can reliably extrapolate from the training data.

- Signal variance σ^2 – This determines the variation of function values from their mean. Small values dictate that the function stays close to its mean, whilst larger values allow for more variation. If the signal variance is too large, the functions will be heavily influenced by outliers.

- Noise variance – The noise variance is used by the GP to allow for noise present in the data.

Through the use of different kernels and the hyperparameters described above, expert knowledge of the workings of BO can be incorporated to account for the expected nature of the underlying objective function. For example, if it is already known that the objective function is a smooth surface then a squared exponential kernel would be suitable and more efficient than any of the Matérn kernels.

Likewise, the hyperparameters can be manually set according to the underlying objective function; if the objective function is very noisy with lots of local optima, then settings the length scale to a low value and the signal and noise variances to high is likely to produce the best outcome. Unfortunately, in order to accurately determine which settings are most suitable, they must be used in the active sampling process and can not be accurately compared post-hoc. As little is known about the objective functions estimated in this thesis, the length scale and signal variance are always optimised. In chapter 3, the noise variance is set according to initial data recorded in the burn-in phase, whereas it is optimised in chapter 5.

2.1.4 Acquisition Functions

The role of the acquisition function is to guide the search for the maximal value of the objective function – the optimum. As acquisition functions are generally defined so that high values of acquisition correspond to potentially high values of the objective function (Brochu, Cora, and Freitas 2010), this is done through maximising the acquisition function itself. By doing this the acquisition function selects the next point at which to evaluate the objective function within each iteration.

A wide variety of acquisition functions exist, however which one is chosen depends on the degree to which the experimental space must be explored or exploited. In this context, exploration refers to sampling in locations where little information is known about the objective function, i.e. where the surrogate variance is large, and exploitation refers to the process of taking advantage of areas in which the objective function is known to be high, i.e. where the surrogate mean is high. The required balance in the exploration-exploitation trade-off varies, however it is gen-

erally desirable to initially explore the space and then move towards exploitation of the areas with a higher surrogate mean. A description of a number of acquisition functions can be found below:

- Probability of improvement (PI) – Whilst this is a commonly used acquisition function, it is highly exploitative and less useful in situations involving high variance, therefore it is not utilised in this thesis.

- Expected improvement (EI) – This acquisition function not only takes into account the probability of improvement, but also the potential magnitude of improvement that a point may make, allowing for a trade-off between exploration and exploitation.

- Upper confidence bound (UCB) – The UCB weights points where there is a high variance to those with a high mean, enabling a trade-off between exploration and exploitation. In its simplest form, this balance between exploration and exploitation is determined by a tuning factor K, such that the GP-UCB at any given point, x, is defined as:

$$UCB(x) = \mu(x) + K\sigma(x) \tag{2.3}$$

2.1.5 Importance Sampling

One problematic possibility in BO is that the algorithm finds one point in the experimental space with a large predicted mean whilst uncertainty at all other points is relatively small, causing the algorithm to repeatedly sample at this location. If such a circumstance occurs when other areas of the space have been sufficiently

explored to be confident that their means are low then this is a satisfactory outcome as it suggests the algorithm has converged on the point at which the objective function is highest. Unfortunately this can also occur before the algorithm has had a chance to explore enough of the surface to be sure of this, and is usually the result of a particularly large peak in the surface. This can result in the algorithm becoming stuck exploiting one location and not exploring the surface.

In order to negate the chance of premature exploitation, the work in chapter 3 employs a technique known as importance sampling. This technique allows for all points in the surface to have a chance of being sampled, however that chance is weighted depending on the value of the acquisition function at each point and raised to the power of a tuning factor.

2.1.6 personalised Bayesian Optimisation

Whilst the BO algorithm that has so far been described is perfectly capable of providing an active sampling method superior to random sampling, it lacks the ability to select a personalised set of parameters for an individual without testing all of the parameter combinations on that individual. As this would be a time-consuming process, in chapters 5 and 6 a personalised Bayesian optimisation (pBO) algorithm is employed to select personalised stimulation parameters. This pBO algorithm is based on work by A. Krause and Ong (2011) who originally used it in a machine learning setting. The pBO algorithm works in the same basic way as the standard one, in that it relies on a GP and acquisition function to select which stimulation parameters to select. The novelty of this algorithm is that it also incorporates a personalised variable as one of its input variables. The personalised

variable is one which is not manipulated to determine it's effect on the objective function, but is a measurable quality of the individual being tested, for example gender, age or IQ. By selecting a personalised variable which accounts for as much of the variability in the objective function as possible, in the case of tES parameter selection, it may be possible to recommend different stimulation parameters to different individuals based on the value of their personalised variable.

2.2 Electroencephalography (EEG)

Electroencephalography (EEG) is a non-invasive neuroimaging method capable of recording the electrical activity present in the brain. This is done by utilising an array of electrodes on the scalp which detect any electrical impulses from the brain and feed the continuous signal back to a recording computer to be analysed. The location of these electrodes is governed by the international 10–20 system, enabling electrode locations to be standardised across different systems. The EEG system used in this thesis was the Starstim 32 device (Neuroelectrics, Spain), capable of recording from up to 32 electrodes. The system has a fixed pair of reference and ground electrodes located on the right mastoid.

2.2.1 EEG pre-processing

Before the EEG can be meaningfully analysed, the data must be pre-processed to remove physiological and motor artifacts such as blinking, eye movement, head movement and participant heartbeat. In this thesis I have utilised the same pre-processing pipeline in all my EEG data. This pipeline is run in Matlab 2017b (Mathworks) using the EEGlab toolbox (Delorme and Makeig 2004) and custom

scripts and is comprised of 3 main steps:

1. **Filter** - filters enable sections of the power spectrum to be filtered out of the data, ensuring that any artifacts within them do not contaminate the sections of interest. The type of filters used in this thesis were linear finite impulse response (FIR) filters available using the EEGlab toolbox. These filters are applied both forwards and backwards to ensure any phase delays induced by the filter are cancelled out. This pre-processing pipeline employed a high-pass filter to the data to remove any linear trends, unless otherwise stated, this was done at 0.1Hz. A low-pass filter was also applied to remove higher frequency noise. This was run at 70HZ unless otherwise stated. In addition, a notch filter was applied at 49–51Hz to remove any line noise from electrical equipment surrounding the EEG system at the time of recording.

2. **Manual artifact rejection** - the time series is manually inspected, and areas in which there are movement artifacts or high levels of linear drift, which were not removed by the filtering, are removed from the data.

3. **Independent component analysis (ICA)** - an ICA is a linear decomposition method which uses a series of spatial filters to transform the data from a time series in which each row of data represents voltages collected from a different recording channel, to a time series in which each row of the data matrix represents the activity present in a component of the data. Each component is maximally independent from any other single component and can be comprised of data sources from multiple recording channels. The components represent information from different sources including cortical and non-cortical sources. In this pipeline, the ICA is run to detect non-cortical sources of noise

such as blinks, eye movement and heart beats. The components representing these sources of noise are identified through visual inspection and removed.

2.2.2 EEG analysis

Power analysis

A commonly used method to analyse EEG data is to decompose the signal into distinct frequency bands including delta (0.5–4 Hz), theta (4–8 Hz), alpha (8–12 Hz), beta (12–30 Hz), and gamma (30–100 Hz). To do this, the signal is transformed into the frequency domain using a Fourier transform, the most common method of doing this being the Fast Fourier Transform (FFT). This algorithm produces the power spectral density (PSD), from which amplitude and phase of any given frequency band can be calculated. All power analysis in this thesis utilised the function 'spectopo' from the EEGlab toolbox (Delorme and Makeig 2004) which uses the function 'pwelch' from the Matlab signal processing toolbox to calculate the PSD using an FFT.

Weighted phase lag index (wPLI)

Weighted phase lag index (wPLI) is a measure of the connectivity between two EEG electrodes. Phase lag-based measures of connectivity account for results due to volume conduction as well as spurious results caused by the measurement of activity from the same source by two electrodes. This is achieved as the phase angle between the two electrodes of interest will have a phase lag of zero when measuring from the same source, indicating volume conduction, which can be avoided. Specifically, phase-lag index measures whether the phase angle is positive or negative, with a

number further from zero indicating higher phase-lag and connectivity.

Chapter 3

Stimulation parameter selection using Bayesian Optimisation

3.1 Introduction

Previous work utilising BO in stimulation parameter selection aimed to identify the optimal phase and frequency of tACS needed to induce retinal phosphenes (Lorenz, Simmons, et al. 2019). In both of the two experiments in the study, two blocks of tACS were administered to the participant, each lasting 5 seconds and consisting of different phase-frequency combinations. In this setup, the input to the BO algorithm consisted of participant's binary choices indicating which of the two blocks elicited more intense phosphenes, and the algorithm then selected the next two combinations to be compared. Over the course of 20 iterations, the BO algorithm was able to identify stimulation parameters that best elicited phosphenes in each individual. Importantly, had Lorenz et al. routinely tested each parameter combination, it would have taken a maximum of 83 minutes in experiment 1 and 51 minutes in experiment 2, whereas using this BO method, it took 5 minutes to find the optimal parameters for each participant. In terms of using BO to select

stimulation parameters, this work serves as an excellent proof of concept. However, the number of participants in both experiments was only 10 as the BO algorithm was entirely run within each participant. This within-subject design when using BO is an excellent way of determining stimulation parameters specific to the individual. Unfortunately, it limits the applicability of the stimulation parameters to future participants, as the BO procedure must be re-run on any new participants. In contrast, utilising a between-subject design, whilst not achieving the same levels of individualisation as a within-subject design, would allow for an initial study to be performed identifying the optimal stimulation parameters for improving performance on a cognitive task. These optimal parameters could then be used on all future participants, without the need to run the BO algorithm again, making this approach particularly appealing to studies with larger cohorts. This is a particularly important issue in studies examining more complex behaviours with noisier responses, as longer behavioural blocks are required to acquire a reliable measure of the effects of the stimulation. In a within-subject design, this leads to a much longer time required to establish an individual's optimal stimulation parameters, for example a task that requires a 5-minute behavioural block to gain a reliable measure would take almost 1.7 hours to run 20 iterations of the BO, and given the larger levels of noise present in more complex behaviours when compared to phosphene perception, is likely to require more than 20 iterations (Letham et al. 2019). Requiring this large amount of time to identify an individual's stimulation parameters is likely to make it impractical for researchers to utilise a within-subject design of BO to examine more complex behaviours, especially in properly powered studies. An additional consideration it that expanding the use of BO to a between-subject design is imperative in order to allow for its future translation to medical

conditions that would be unsuitable for a within-subject setting.

In this chapter I utilise a between-subject BO design to target a cognitive ability of key importance in human behaviour: the transfer of abilities and knowledge between different tasks. Transfer is crucial to human intelligence as it allows for the adaptation to new tasks and increased flexibility in learning. Transfer between tasks can be divided into near and far transfer, referring to how similar the two tasks are, with near transfer being between similar but not identical tasks, and far transfer being between dissimilar tasks which share a general skill (Taatgen 2013). Specifically, I targeted near transfer from tasks in which the instructions were well practised to those involving combinations of instructions new to an individual. In addition to the reasons previously stated, a between-subject design was employed because the task involved learning, and therefore a repeated measure of the participant's performance under different conditions is undesirable.

The behavioural paradigm employed in this chapter investigated rapid instructed task learning (RITL): the ability to rapidly learn a new task from instructions alone (Cole, Bagic, et al. 2010). Previous research by Cole et al. (2010) recorded fMRI and MEG whilst participants completed a RITL task and found that the processing of novel tasks relies on rule integration in the dorsolateral prefrontal cortex (DLPFC) before higher level integration in the anterior prefrontal cortex (aPFC). In contrast, learned tasks involve a reversal of this activity pattern. These neuroimaging results highlight the crucial role the DLPFC plays in RITL, and suggest it may be an ideal target for tES. Additionally, previous research has suggested that tRNS may be beneficial for improving the transfer of abilities from learnt to novel tasks (Snowball et al. 2013; Popescu et al. 2016; Pasqualotto 2016; Brem et al. 2018), making it a promising tool for modulating transfer in this behavioural

paradigm. Specifically, in this chapter BO was employed to determine the parameters of transcranial random noise stimulation (tRNS) over the dorsolateral prefrontal cortex (DLPFC) most able to improve transfer of task-general abilities from previously learnt tasks to new task combinations.

3.2 Methods

Ninety-five participants were recruited for this study. Participants were excluded from the study if they failed to attend the sessions on consecutive days or failed to achieve an accuracy of 80% in the second half of the first session; this was to ensure participants learnt the rule combinations presented to them. Fourteen participants failed to achieve the required accuracy of 80% in the second half of the first day of testing (for similar criterion see Cole, Patrick, et al. 2018). One participant did not attend the second testing day. Therefore, the data from these participants was not included in the optimisation. Eighty participants' data points were used in the optimisation (53 female, mean age±standard deviation (SD) = 22.4 ± 5.91) including 35 in the burn-in phase.

Participants were screened for exclusion criteria for tES including a personal or family history of epilepsy, neurological condition, head trauma and medications (see Appendix A.1 for full list). After providing informed consent, participants completed the matrix reasoning component of the Wechsler Abbreviated Scale of Intelligence - Second Edition (WASI-II). This measure was taken in order to later control for variation in intelligence between participants. It should be noted that due to laboratory setup this test was conducted with the experimenter seated behind the participant in contradiction to the test instructions. The only caveat with

such a setup is that the score cannot be compared reliably to the norms in this test, but they could be used reliably to capture variance in the performance within our sample, which was the motivation for administering this task. The 32-channel EEG system was then set up and two three-minute resting state recordings were captured with participant eyes open and closed. The experimental task was then explained to participants, with each participant completing the same four example trials, before starting the full task.

3.2.1 Behavioural Task

The behavioural task used in this experiment was the permuted rule operations cognitive paradigm developed by Cole et al. (2010) to measure RITL. In this task participants were sequentially presented with a number of three-part instructions composed of a logic, semantic and motor rule. Participants then had to apply the memorised rules to three pairs of word stimuli that were sequentially presented to them and respond 'true' or 'false' depending on whether the stimuli obeyed the rule or not. In order to register a valid response, participants had to respond within 3 seconds of seeing the second word of each pair (Figure 3.1).

The RITL task allows for 64 combinations of 12 unique rules (4 logic, 4 semantic, 4 motor). During the first day participants were allocated 4 combinations of rules to learn whilst receiving stimulation. The rule combinations were selected randomly whilst ensuring each participant was exposed to every possible rule in the first day of testing, but the exact stimuli on the screen differ. Participants completed 120 trials in total (30 different trials x 4 rule combinations). On the second day participants completed 60 trials of the rule combinations from day one (15 different trials x 4

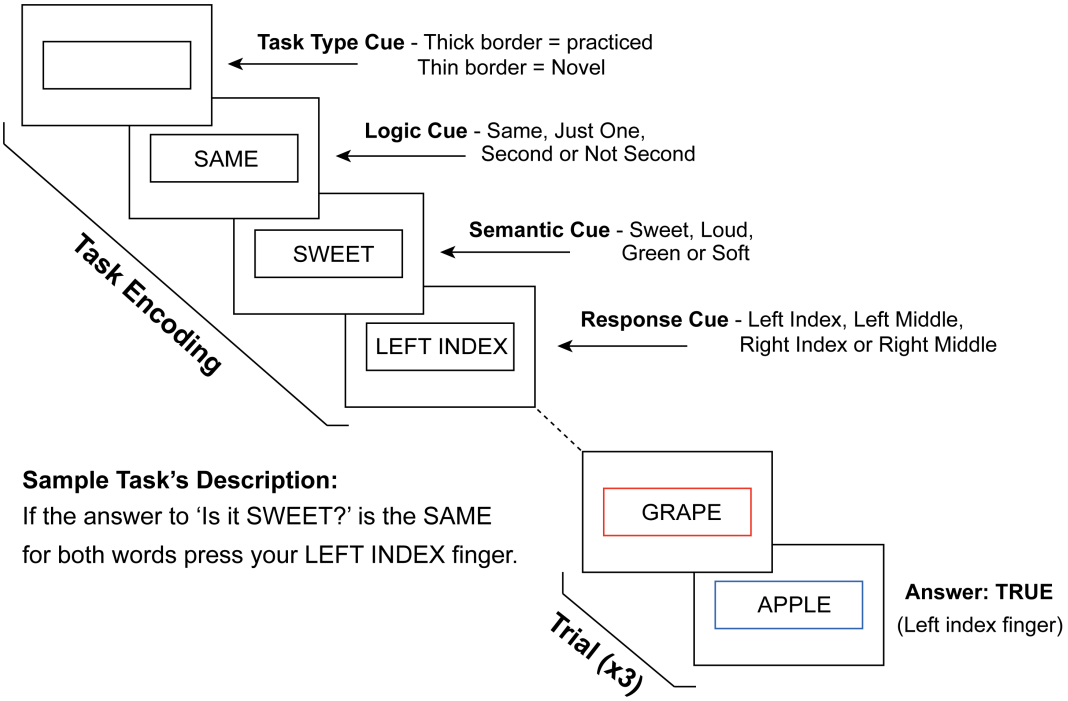


Figure 3.1: **A summary of the RITL task.** Participants were presented with 120 sets of instructions, each followed by three pairs of words. For each pair of words they responded true or false as quickly and as accurately as possible. In this example participants were sequentially presented with the instructions SAME, SWEET, LEFT INDEX. They were then presented with the word pair GRAPE and APPLE, to which they would respond by pressing the button under their LEFT INDEX finger to indicate a true response as both grapes and apples are sweet in taste.

practiced rule combinations) and 60 trials of completely novel rule combinations, randomly interleaved.

3.2.2 Behavioural measures

The presence of both previously learnt and novel rule combinations in session 2 allowed for a measure of transfer to be obtained for each participant. EZ diffusion modelling was used to attain a singular value combining participant accuracy and

reaction times. EZ diffusion modelling allows for the calculation of a measure of how quickly a participant reaches the correct answer in the form of a drift rate value (Wagenmakers, Van Der Maas, and Grasman 2007). The degree of overlap between the rule combinations shown in the second day with those in the first day, defined the degree of transfer. Specifically, the trials were coded as a transfer distance of 0 for rule combinations in day 2 that were previously learnt in day 1. Trials were coded as a transfer distance of 1 for rule combinations that included two rules previously seen together in day 1, while trials with combinations that were not seen together in day 1 were coded as a transfer distance of 2. The gradient of the drift rates across the three transfer distances (0–2) was then calculated for each participant and used as a measure of performance, with more positive gradients interpreted as indicating better transfer, similar to a previous published study (Cole, Reynolds, et al. 2013).

3.2.3 Stimulation and EEG

Stimulation was delivered, and EEG was recorded using the Starstim 32 device (Neuroelectrics, Spain) with electrodes set up according to the standard 10-20 system. Participants received tRNS on day one of testing. Stimulation was delivered at locations F3 and F4 for 20 minutes plus 30 seconds ramp up/down in the active condition, and for 30 seconds ramp up/down only for the sham condition. tRNS was applied at the beginning of the behavioural task. The behavioural task lasted approximately 40 minutes, allowing for the investigation of both online and offline effects in session 1. During the offline period of day 1 and during day 2, participants' EEG was recorded. EEG data is not reported in this chapter as it is not

considered relevant to the present research question.

After a burn-in period of 35 participants, stimulation current and frequency was determined through BO for the remaining 45 participants. Stimulation was delivered between 0.4–1.9mA (peak-to-peak) at intervals of 0.1mA and 0–500Hz in bands of 100Hz. The parameter combinations chosen for the burn-in phase were selected to provide a relatively even spread of data across the experimental space and included 5 participants who received sham tRNS in order to determine an estimate of the level of noise to expect at each behavioural response.

3.2.4 Bayesian Optimisation

After the burn-in phase of 35 participants was complete, stimulation parameters for all subsequent participants were determined via BO whilst maintaining a double-blind paradigm. Based on the current and frequency of stimulation received, as well as each previous participant's drift rate gradients, the algorithm selected new stimulation parameters for each participant whilst attempting to find the parameter combinations that led to the maximum behavioural response.

The BO algorithm was an adapted version of a previously published algorithm (Hadida et al. 2018). All data was normalised using an error function, after which a GP was trained on the data using a matern 32 kernel and an initial sigma set to the standard deviation recorded from the scores of the 5 participants who received sham in the burn-in phase. This sigma value was optimised to appropriately fit the data. Once the GP was trained, an upper-confidence bound (UCB) acquisition function was utilised to determine the attractiveness of sampling at each possible point in the experimental surface. These possible sampling points were constrained

to every 0.1mA along the current dimension and to five 100Hz frequency bands ranging from 0–500Hz. To avoid the algorithm getting stuck at local maxima and to encourage exploration, importance sampling was implemented when selecting the next points to sample. This involved linearly re-scaling all UCB scores between 0 and 1, with 1 being more optimal, and then raising them to the power of 2, therefore making more attractive locations more likely to be sampled whilst still allowing less attractive locations to be sampled on occasion. Due to logistical constraints, the optimisation algorithm was run after every 2–3 participants depending on scheduling, where upon stimulation parameters were determined for the next 2–3 participants. Determining the stimulation parameters of 2-3 participants at one time enabled multiple participants to be tested at the same time in the two-day paradigm and enabled testing to progress faster. The protocol for this study was preregistered on AsPredicted (http://aspredicted.org/blind.php?x=jw9t7w).

3.3 Results

Analysis of the BO algorithm demonstrated that it was able to improve the best predicted performance found from the fixed search burn-in employed on the initial 35 participants (Figure 3.2). Additionally, the average uncertainty estimated by the GP was significantly smaller when utilising BO compared to fixed search (BF=1,093,000, $t(69)=6.56$, $p<.001$, $d=1.63$).

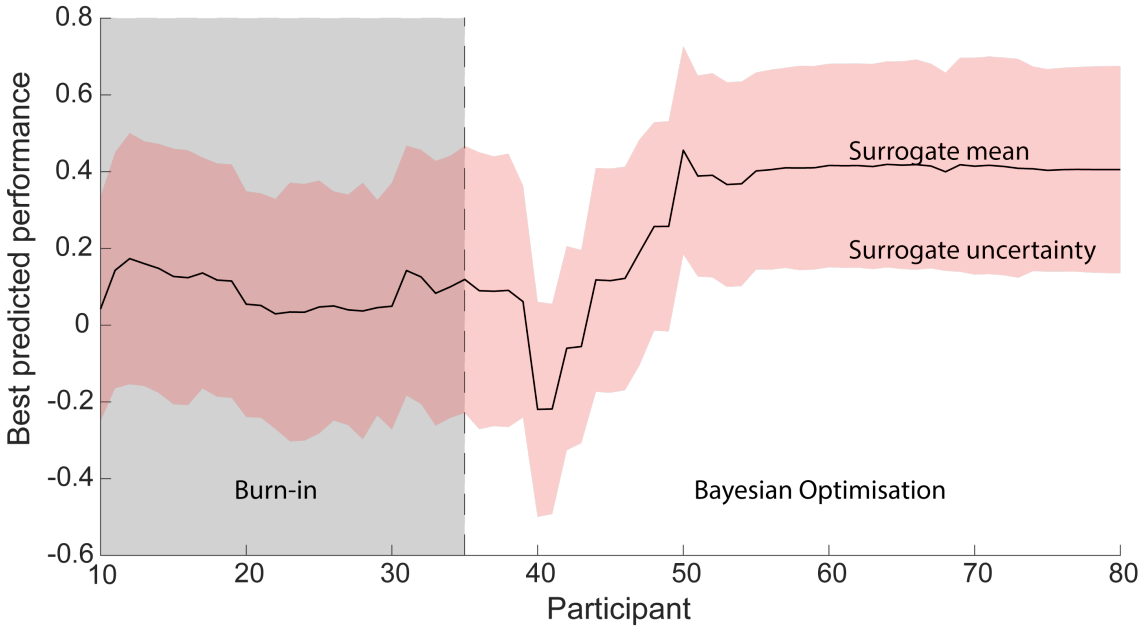


Figure 3.2: **Highest performance predicted by the Gaussian process surrogate mean after each participant**. The surrogate uncertainty is shown by the shaded area. The first 35 participants were sampled using fixed search to determine their stimulation parameters, all participants after that had stimulation parameters chosen by the BO algorithm.

Figure 3.3 shows the drift rate gradient predicted by the GP model at each stimulation parameter combination based on data from all 80 participants. A clear peak was found at 0.8mA, covering the 100–200 and 200–300Hz frequency bins, with the highest point at 200–300Hz.

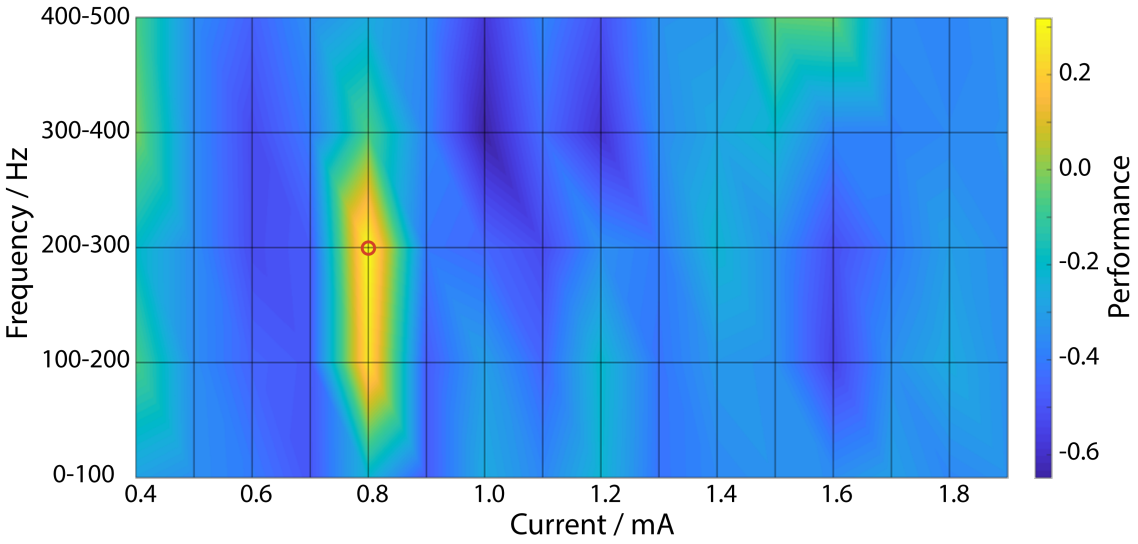


Figure 3.3: **Performance predicted by the Gaussian process model at each stimulation parameter combination after 80 participants.** Performance is indicated by the surface colour, demonstrating a peak at 0.8mA over the 100–200 and 200–300Hz frequency bins. The 200-300Hz bin is the highest, denoted by the red circle. Performance relates to the level of transfer of abilities from rule combinations learnt in day 1, to novel rule combinations seen in day 2 of the paradigm.

3.3.1 Simulation analysis

Simulation analysis were run to determine the ability of the BO algorithm to detect a peak in simulated data given different effect sizes of stimulation. To do this, different effect sizes were simulated by manipulating the level of noise introduced into a simulated surface. The underlying surface was a completely flat 16x5 surface centred around 0, with one clear peak, peaking at 1. When the BO algorithm sampled the surface, noise was added to the sampled value. By manipulating the standard deviation of the added noise, different effect sizes were simulated between the peak and the rest of the flat surface.

To keep the simulation as true to the experiment as possible, the simulation

ran for the same number of iterations as the real experiment (80). It also had a near-identical burn-in period, with the five sham participants removed as 0mA was not included in the search space. Both the BO algorithm and a random search algorithm were run on the surface to demonstrate the difference between the two. The algorithms were run 100 times at each simulated effect size, with performance measured as the Euclidean distance between the true optimum and the one selected by the BO or random search algorithm after the total number of iterations.

Results of the simulation showed that the BO algorithm achieved a smaller Euclidean distance between the selected optimum and the true optimum than the random sampling algorithm. This result suggests the BO algorithm is more successful at identifying optima than the random algorithm at a broad range of effect sizes. Interestingly, the simulation showed that very large effect sizes (d>3) were needed to enable the BO to reliably identify an optimum within 1 space of the true optimum (Figure 3.4).

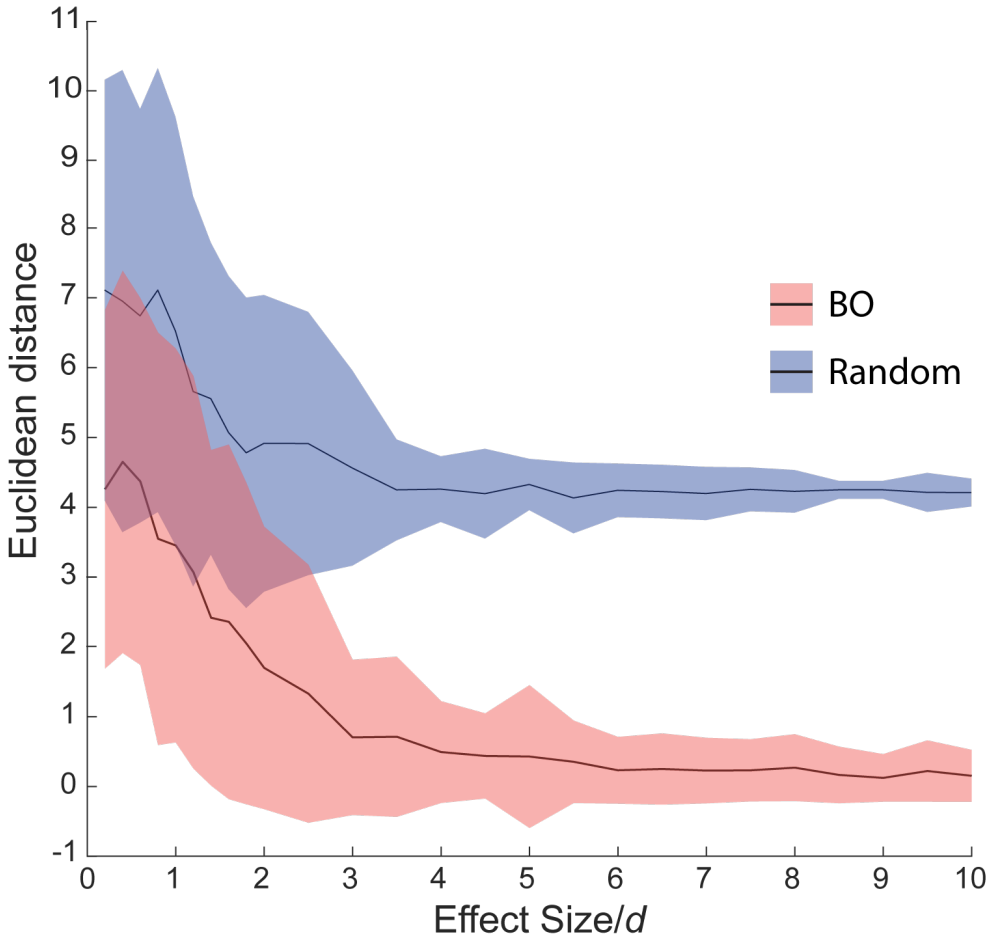


Figure 3.4: **Results of simulations examining the ability of the Bayesian optimisation and random search algorithms to find the optimum at different simulated effect sizes.** Effect sizes were simulated for a single peak in a flat surface by introducing different levels of noise whilst keeping the height of the peak constant. Performance of the Bayesian optimisation and random search algorithms was measured as the Euclidean distance between the true optimum and the one selected by the algorithms. The figure shows mean and standard deviation of Euclidean distance.

3.4 Discussion

In this chapter I examined the suitability of BO to find the optimal tRNS parameters to promote transfer in the RITL task. Crucially this was carried out using a between-subjects design, in which each stimulation parameter selected for testing was used on a different individual, allowing for the identification of non-subject specific stimulation parameters. The algorithm identified the optimal tRNS parameters as 0.8mA and 200-300Hz for improving transfer (Figure 3.3). These results indicate that the use of BO in a between-subject design to identify stimulation parameters is feasible and relatively simple to run. Further, this methodology could provide a data-driven method for parameter selection that can not only be utilised in tES, but in a variety of neurointerventions.

These results could not only be meaningful for the future use of BO in brain stimulation, but support underlying mechanistic theories of tRNS. Previous work using tACS demonstrated that cortical excitability is modulated using 140Hz tACS, but not 80Hz (Terney et al. 2008), similar results have also been found using tRNS across these ripple frequencies (Moliadze, Antal, and Paulus 2010) and have suggested that 'ripples' between 140–220Hz are involved in learning and memory (Ego-Stengel and Wilson 2010; Girardeau et al. 2009). Whilst the results of this chapter identify the optimal frequency bin as 200–300Hz, 100–200Hz was also predicted to improve performance. The combination of these frequency bins would encompass the 140–220Hz band predicted by the ripple oscillation theory, potentially supporting it as an underlying mechanism of tRNS. However, further work needs to be done to determine whether this is a reliable effect. Additionally, the use of a narrow frequency band could have lead to the stimulation not having the same effects

on cortical activity as a classical broad frequency band would have, potentially limiting the stimulation's effects.

Detailed visualisation of the performance of the BO algorithm, shows that it was able to improve upon the best found parameter from the burn-in phase of the initial 35 participants (Figure 3.2), which tested a series of fixed parameter combinations spread evenly throughout the search space. This burn-in phase provided a basic understanding of the search space, allowing the BO algorithm to then exploit areas in which performance was high, and further explore areas in which variance was high. This is illustrated by the drop seen in best predicted performance in Figure 3.2 at around 40 participants. Here, the BO algorithm re-tested the best parameters found in the burn-in phase, and found they were not as beneficial as originally thought. The algorithm then moved on to test other more successful parameters, signified by the increase in predicted performance between participants 40–50. These best parameters were then retested, as well as other parameters explored, indicated by the plateau in best predicted performance from around participant 50. Behaviour such as this is exactly what is expected of an active machine learning technique such as BO.

The ability to balance the exploitation of potentially beneficial parameter combinations and the exploration of unknown ones is key in developing an efficient optimisation algorithm. This is demonstrated by the simulation analysis carried out in this chapter, which showed the superior performance of the BO algorithm against a random search algorithm. These results also highlighted the impact of variation in behavioural responses on the effectiveness of BO in human cognition, as the simulation demonstrated that a Cohen's d $\geq$3 would be needed to allow the estimated optima to be within 1 square of the true optima. Whilst the increased performance

compared to random search is a positive, the requirement for such a large effect size in the BO algorithm calls into question the reliability of the results in this chapter, as both behavioural measures such as transfer and brain stimulation experiments can have large inter- and intra-individual variability, as discussed in chapter 1. This variability could introduce noise into the BO environment, and may have meant that the BO algorithm was unable to reliably identify a peak in the parameter space and instead modelled noise in the data. In computer science, BO is commonly applied in situations in which there is low/moderate levels of noise, a large amount of data, or both (Joy et al. 2016). Previous studies in low-noise situations have suggested a sample size of 10D or 20D, where D is the number of dimensions present in the algorithm (Snoek, Larochelle, and Adams 2012; Hernández-Lobato, Hoffman, and Ghahramani 2014; Nguyen et al. 2017). More recent work employing BO in a noisy environment utilised 25D (Letham et al. 2019); due to the high levels of anticipated variability 40D was used in this study, however this may still not be enough. In the present study, and typically in most human behavioural studies, there was a large number of uncontrolled factors which could introduce noise into the collected data (for review see Ridding and Ziemann 2010). These inter– and intra–individual factors, such as gender (Chaieb, Antal, and Paulus 2008; Fumagalli et al. 2010; Russell et al. 2017), age (Bartrés-Faz, Rossi, and Pascual-Leone 2016), genetics (Antal, Chaieb, et al. 2014; Teo et al. 2014), electrode location and placement (Wörsching et al. 2018), stimulation intensity (Ammann, Lindquist, and Celnik 2017; Chew, Ho, and Loo 2015), operator training and insufficient protocol reporting (Bikson, Brunoni, et al. 2018). Additionally stimulation effects have been shown to be brain-state dependent (Cohen Kadosh, Soskic, et al. 2010; Feurra et al. 2013; Neuling, Rach, and Herrmann 2013), a factor that whilst influenced through

the behavioural task being completed, can never be fully controlled. The total contribution of these factors to the noise levels in this experiment is what is modelled by the GP as noise. Although it should be noted that the underlying assumption of the GP is that this noise is Gaussian, which could be a limitation of this method, in that not all of the underlying sources of noise will be Gaussian, for example reaction times can not be lower than 0ms therefore limiting the distribution of this variable. It should be noted, that previous work conducted similar simulations to the one in this chapter and showed that the BO algorithm was able to find an optimum reliably close to the true one at effect sizes of 0.3 (Cohen's d). This study used a larger simulated experimental surface of 361 points compared to the 80 used in the simulation in this chapter and had a smoother transition between the peak and the rest of the surface, allowing for easier modelling by the GP. However this is still a surprising difference in results and whilst it could partly be explained by the use of a different acquisition function, it calls for a further validation of the results of the current chapter.

The choice to use drift rate as an integral part of the measure of participant's behavioural responses was a decision that could have a significant impact on the success of the BO algorithm. Drift rate was chosen because it provided a single measurement of performance combining both accuracy and reaction time, which is suitable for BO as it requires a single measure to optimise. However, drift rate is only one of the three components of the EZ–diffusion model, with the other two being boundary separation and non-decision time. It is possible that one of these alternative EZ–diffusion components could better capture the change in behaviour in response to tES, or alternatively that the EZ–diffusion model is not the most appropriate model to capture such changes. However, as the EZ–diffusion model

is a well tested cognitive model which combines accuracy and reaction time into one metric in the form of drift rate, it was a sensible one to choose in the current experiment involving BO.

In summary, in this chapter I presented the first use of BO to determine parameters of tES in a complex cognitive task. In identifying the optimal stimulation parameters to improve transfer in RITL I found promising results as the algorithm identified an optimal stimulation current intensity and frequency. However, simulations left doubt about these results and further validation is required. This validation study will be presented in chapter 4.

Chapter 4

Validating parameter selection with Bayesian Optimisation

4.1 Introduction

Previous work utilising BO in stimulation parameter selection ran the BO algorithm independently for each participant (Lorenz, Simmons, et al. 2019). Similar stimulation parameters were found across participants, adding validity to the results. However, in some situations, such as when using BO in more complex cognitive behaviours or using an experimental design incorporating multiple participants, it is desirable to fully validate any results on a separate sample of participants. In this chapter, I aim to validate the results from chapter 3 by testing the effects of the BO derived stimulation parameters against sham stimulation. The experimental setup and the behavioural task were identical to chapter 3, however, in a double-blind protocol, participants were divided into 2 independent groups and either received active or sham tRNS in session 1.

4.2 Methods

The behavioural task utilised in this chapter was identical to that in chapter 3. Participants learnt 4 rule combinations of the RITL task in the first session and then returned for the second session the next day and completed the task comprised of both learnt and novel rule combinations. In this protocol, tRNS was delivered during the first session as either sham or active stimulation. In the active tRNS group, the stimulation parameters used were those determined using BO in chapter 3 (i.e. 0.8mA and 200-300Hz). This study was pre-registered on AsPredicted (http://aspredicted.org/blind.php?x=7h843d).

4.2.1 Participants

The required sample size was calculated using the Bayes Factor Design Analysis ($BF_{10}DA$) package in R (Schönbrodt and Stefan 2018). To do this, alternative and null hypotheses were calculated for a one-tailed between-subjects t-test using 10,000 simulated studies, with an effect size of 0.6 for H1; all other settings were left as standard. The simulation demonstrated that 70 participants were required to achieve a power of 80% with the BF_{10} set at 6. In order to allow a more conservative estimate of the effect size in this study, 80 naive participants were tested across the two groups. However, to avoid testing more participants than was necessary a sequential sampling technique was employed in which a Bayesian t-test was run after participant was tested. Testing would be stopped if the Bayes Factor of the t-tests reached 6 or 1/6 as this would indicate sufficient evidence for either the alternative or null hypothesis.

Participants were screened for exclusion criteria for tES including a personal

or family history of epilepsy, neurological condition, head trauma and medications (see Appendix A.1 for full list). All participants gave informed consent. In total, 93 participants were tested, 12 of whom did not score highly enough (80% accuracy) in day 1 to participate in day 2 and one of whom did not complete both sessions. Eighty participants completed the experiment (37 females, mean age±SD = 22.35±6.86). During session 1, these participants were allocated to either the sham or active stimulation groups, whilst maintaining a double-blind protocol. In order to match the two groups as much as possible, a minimisation procedure was applied which assigned participants to the two groups. The procedure aimed to minimise the between-group differences in the distribution of the selected features, therefore providing groups that were as equally matched as possible. In this case, the procedure matched participants on gender, age and IQ (Treasure and MacRae 1998).

4.2.2 Stimulation

As in chapter 3, participants received tRNS on day 1 of testing. Stimulation was delivered using a Starstim 32 device (Neuroelectrics, Spain) at locations F3 and F4. Stimulation current and frequency was determined by allocation to either the active or sham stimulation groups. Active stimulation lasted for 20 minutes plus 30 seconds ramp up/down and stimulation was delivered at a current of 0.8mA peak-to-peak and stimulation frequencies were filtered between 200-300Hz. Sham stimulation used all the same parameters as active stimulation but only involved the ramp up, immediately followed by ramp down of the stimulation.

4.2.3 EEG processing

All EEG data was pre-processed in Matlab 2017b (Mathworks) using a combination of EEGLAB (Delorme and Makeig 2004) and custom code. The pro-processing pipeline was the same as that described in chapter 2. EEG data was filtered between 0.1 and 70, after which it was manually inspected and any areas including movement artifacts, or noisy channels were deleted. An ICA was performed and components representing heart beats and ocular artifacts such as blinks and eye movements were removed from the data. For the presentation of stimuli, EEG data was epoched to -0.5–2.5 seconds relative to the onset of the first stimuli within each trial and baselined to the 0.5 seconds before the stimuli onset. For the response periods, the EEG data was epoched to -3.5–0.5 relative to the response and baselined to 0.5 seconds after the response.

4.3 Results

There was no difference between active and sham tRNS groups in respect to age (BF_{10}=0.232, t(78)=0.011, p=0.991, d=0.003), performance IQ (BF_{10}=0.278, t(78)=0.641, p=0.524, d=0.143) or gender (BF_{10}=0.272, X^2(1,n=80)<0.001, p=0.987, Phi=0.002).

4.3.1 Group effects of stimulation

The sequential Bayesian t-tests demonstrated no effect of tRNS on transfer when compared to sham. The final Bayes Factor was 0.301 suggesting moderate evidence for the null hypothesis (Figure 4.1). To test for any differences in variance between

the groups, which would be relevant to the performance of the BO algorithm, Levene's test of equality of variance was performed and no difference was found between the active and sham groups (F=2.847, p=0.096).

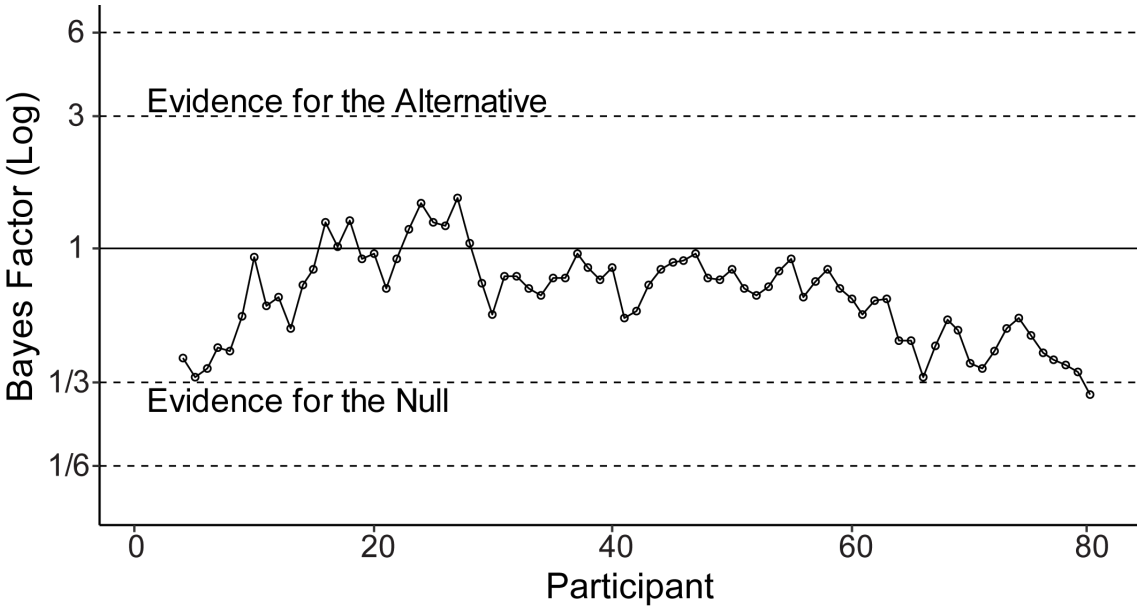


Figure 4.1: **Results of the one-tailed sequential Bayesian t-tests** testing whether stimulation had a positive effect on participant performance in the RITL compared to sham stimulation. The final Bayes Factor was 0.301, indicating moderate evidence for the null hypothesis.

Further exploratory analyses aimed to determine whether there was a stimulation effect on the individual drift rates in day one and day two. Drift rate was calculated separately for the first half and last half of the task completed in day 1. A repeated measures ANOVA was performed with time and stimulation group as factors. This analysis found an effect of time on drift rate (BF_{10}=3.964x10^{11}, F(1,78)=93.156, p<.001, η_p^2=0.544) demonstrating that participant drift rate increased with practice, but no effect of stimulation group (BF_{10}=0.269, F(1,78)=0.038, p=0.845, η_p^2<.001) or any interaction between group and time was found (BF_{10}=0.807, F(1,78)=2.873, p=0.094, η_p^2=0.036), or main effect of stimulation group .

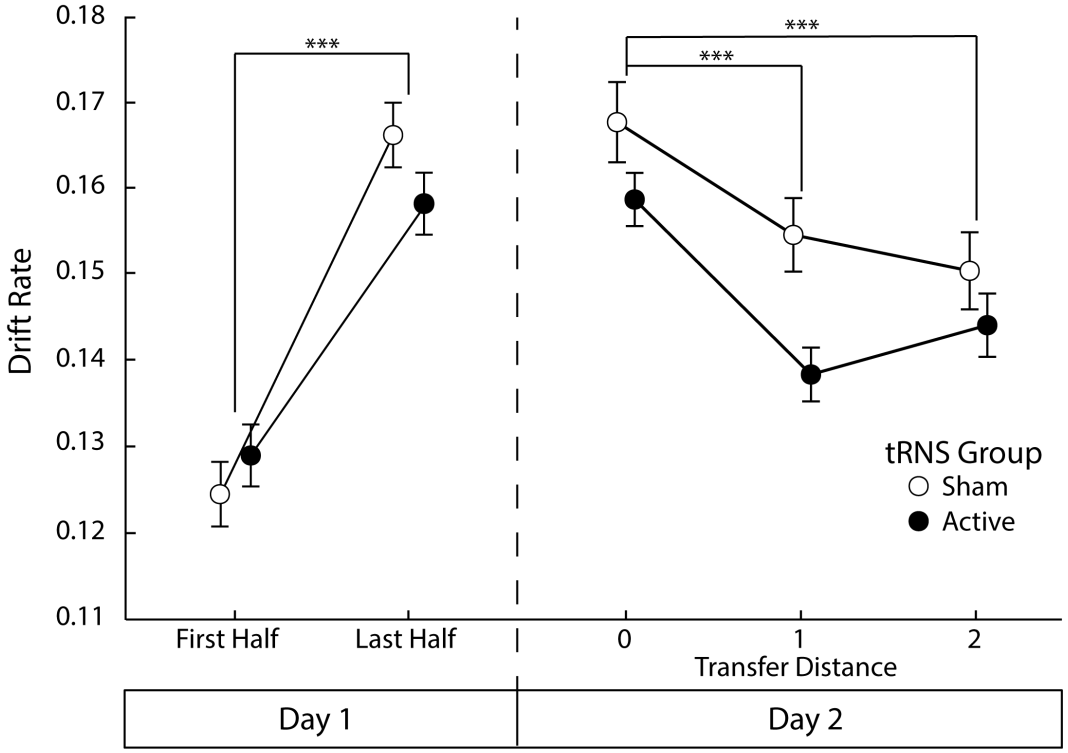


Figure 4.2: **Drift rates during both days of the experiment.** Day 1 drift rates are shown for the first and second half of the session. Day 2 drift rates are shown for the three different transfer differences across the whole of the second session.Error bars denote the standard error. ***<.001.

A repeated measures ANOVA was performed to determine whether there was any effect of stimulation on the drift rates of the three different transfer distances presented in day 2 (Figure 4.2). An effect of transfer distance was found (BF_{10}=955.499, F(2,156)=11.660, p<.001, η_p^2=0.130), but no effect of stimulation group (BF_{10}=0.542, F(1,78)=0.766, p=0.384, η_p^2=0.010), or any interaction (BF_{10}=0.173, F(2,156)=0.861, p=0.425, η_p^2=0.011). Bonferroni-corrected post-hoc tests demonstrated that the effect of transfer distance was due to a decrease in in drift rate in trial with a transfer distance of 1 compared to 0 (t=4.512, d=0.504, p<.001) and 2 compared to 0 (t=3.911, d=0.437, p<.001), whilst there was no

difference between transfer distance of 1 and 2 (t=0.230, d=0.026, p=1.00).

A linear regression was performed to investigate participants' demographic effects on objective function. The regression equation was not found to be significant (F(3,76)=0.540, p=0.656) and no effect was found of gender (BF_{10}=0.304, β=0.078, t=0.669, p=0.505), IQ (BF_{10}=0.378, β=-0.098, t=-0.834, 0.407), or age (BF_{10}=0.253, β=0.044, t=0.375, p=0.708).

4.3.2 Comparison of validation and optimisation data

Data collected from participants who received stimulation at 0.8mA and 200-300Hz was combined with all data from chapter 3. A GP was trained on this data, which demonstrated that there was no increase in performance at this stimulation parameter (Figure 4.3). Participant demographic data was not found to differ between the BO-experiment in chapter 3 and the validation experiment in chapter 4, for participant IQ (BF_{10}=0.220, t(158)=-0.741, p=0.460, d=-0.117), age (BF_{10}=0.171, t(158)=0.099, p=0.921, d=0.016), whether English was their first language (BF_{10}=0.174, t(158)=0.199, p=0.843, d=0.031), or the number of presented words that they did not know the meaning of (BF_{10}=0.172, t(158)=0.118, p=0.906, d=0.019).

4.3.3 Blinding Analysis

Analysis of the effectiveness of the blinding was carried out to determine whether there was a difference between stimulation and sham conditions in terms of whether participants thought they had received stimulation. In the active stimulation group, 11 participants thought they had received active tRNS, whilst 7 participants in

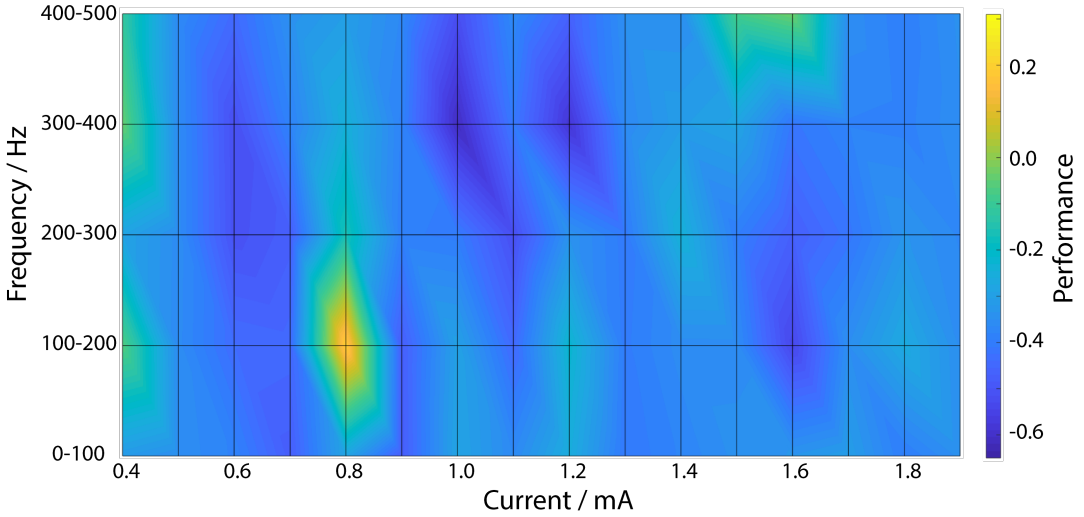


Figure 4.3: **Performance predicted by the GP model at each stimulation parameter combination.** The GP was trained on data from 80 participants in chapter 3, as well as 40 participants tested at 0.8mA and 200-300Hz in chapter 4. The peak previously present in chapter 3 at 0.8mA and 200-300Hz has reduced.

the sham group thought they had received active tRNS. A Chi-squared test revealed no difference between the groups (BF_{10}=0.495, X^2(1,n=80)=0.904, p=0.342, Phi=0.106), suggesting blinding was effective.

Additional analysis investigated the sensations reported by participants in the questionnaire as a result of stimulation (see Appendix A.2 for questionnaire). No differences were found for ratings of itching (BF_{10}=0.073, X^2(2,n=80)=0.,749, p=0.788, V=0.097), pain (BF_{10}=1.547, X^2(1,n=80)=5.073, p=0.168, V=0.252), burning (BF_{10}=0.136, X^2(2,n=80)=4.966, p=0.291, V=0.249), warmth (BF_{10}=0.062, X^2(2,n=80)=0.501, p=0.778, V=0.079), metallic taste (BF_{10}=0.099, X^2(1,n=80)=1.065, p=0.705, V=0.115) or fatigue (BF_{10}=0.102, X^2(3,n=80)=1.313, p=0.788, V=0.128).

4.3.4 EEG analysis

Day 1 resting state EEG data

rs-EEG data from 78 participants (38 sham) were used in this analysis as rs-EEG data from two participants was rejected due to high levels of artefact noise. In order to test whether day 1 rs-EEG values predicted transfer in day 2, separate linear regressions were run using theta power in F3/F4/P3/P4 and wPLI values between F3/F4, F3/P3 and F4/P4 as predictors. Additionally, the theta/beta ratio was extracted from Cz and used to predict transfer. None of these variables were found to be significant predictors of transfer (all $p>0.1$).

Day 2 task EEG data

EEG data from 66 participants (34 sham) was used in the EEG analysis, as three participant's data were rejected due to high levels of noise in the data, whilst an error was found in 11 participant's EEG data in which the event triggers were not reliably recorded in the data, rendering the data unreliable, leading to this data also being rejected.

Theta power was measured separately for the three different transfer distances (0, 1 and 2) in the DLPFC (F3/F4 electrodes) and PPC (P3/P4 electrodes) during stimuli presentation and participant response. A repeated measures ANOVA with the factors transfer distance, hemisphere (F3/F4 or P3/P4) and stimulation group was performed. Analysis of the presentation period showed a main effect of transfer distance in parietal electrodes ($F(2,128)=6.481$, $p=.024$, $\eta_p^2=.092$), but no effect of hemisphere ($F(1,64)<0.001$, $p=0.995$, $\eta_p^2<.001$), stimulation group ($F(1,64)=1.416$, $p=.571$, $\eta_p^2=.022$), or any interactions. Bonferroni-corrected post-hoc tests showed

that the effect of transfer distance was due to a decrease in parietal theta in trials with a transfer distance of 1 compared to 0 (t=2.751, d=0.339, p=0.023) 2 compared to 0 (t=3.296, d=0.406, p=0.005), whilst there was no difference between transfer distances of 1 and 2 (t=0.273, d=0.034, p=1.00). No main effects or interactions were found in the DLPFC after correction for multiple comparisons.

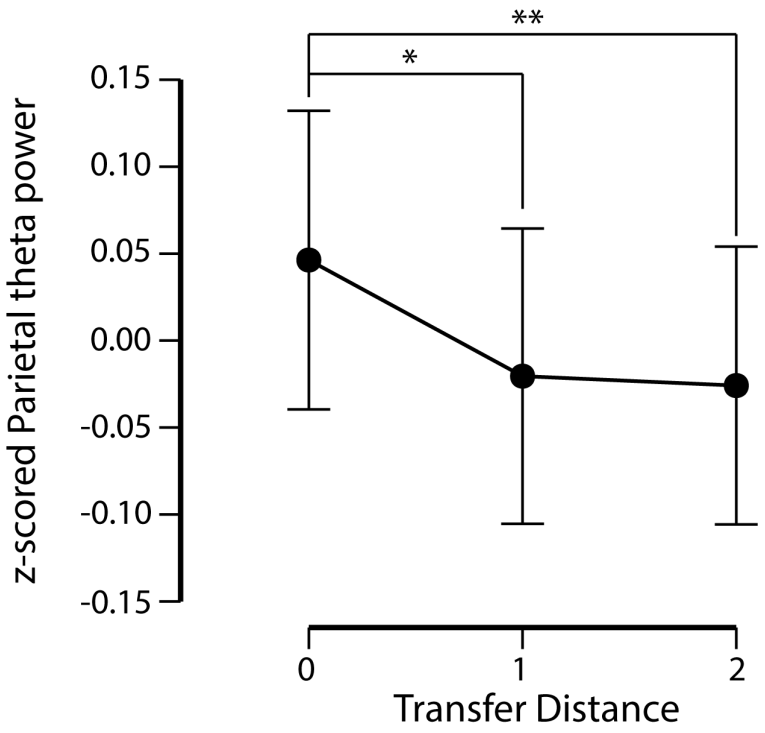


Figure 4.4: **Parietal theta power at different transfer distances** showing an increase in parietal theta in transfer distances of 0 compared to 1 and 2. error brs denote standard error, <0.05 (*), <0.01 (**), and <0.001 (***)

Similar results were found during the response period, with a main effect of trial type in parietal electrodes (F(2,128)=5.883, p=0.024, η_p^2=0.084), but no effect of hemisphere(F(1,64)<0.001, p=0.995, η_p^2<.001), stimulation group (F(1,64)=0.837, p=0.728, η_p^2=0.013), or any interactions. As in the presentation period, Bonferroni corrected post-hoc tests showed that the effect of trial type was due to a increase

in parietal theta in learnt trials comparison to near (t=2.855, d=0.351, p=0.017) and far trials (t=2.628, d=0.323, p=0.032), whilst there was no difference between near and far trials (t=0.300, d=0.037, p=1.00). No main effects or interactions were found in the DLPFC after correction for multiple comparisons.

Connectivity measures in the form of wPLI values were extracted for the presentation and response phases of the task; however after corrections for multiple comparisons, no effects or interactions were found when performing the same repeated measures ANOVAs on the wPLI values. To further investigate the role of DLPFC and PPC theta power in RITL, participant day 2 drift rate was predicted using a regression model with the predictors theta power, trial type and group. However, no effects survived multiple comparison corrections.

4.4 Discussion

In this chapter, I aimed to validate the results of chapter 3, which identified the 'optimal' stimulation parameters for improving transfer in the RITL task using BO. To achieve this, in this chapter I compared the effects of the identified stimulation parameters to sham stimulation. The results show that utilising the BO-derived stimulation parameters did not improve transfer. The results clearly indicate that whilst the algorithm employed in chapter 3 may have detected a positive effect of tRNS at 0.8mA and 200–300Hz, this effect, if reliable, yielded stronger effects than other parameters, but not sham tRNS.

The simplest explanation of these results is that there is no effect of stimulation on transfer in this task. Whilst this is a valid explanation, it is possible that the effects of tRNS truly existed in the participants tested using the optimum stimula-

tion parameters in chapter 3, but were not effective when applied in a larger, more homogeneous sample in chapter 4. Indeed, previous unpublished work from this lab has shown that tRNS delivered at 1mA over the dLPFC can improve performance in the RITL task. Whilst the additional use of the parameters from this previous study would have been ideal in this validation study, it would have required a large sample size and therefore was not feasible. However it should be noted that this previous work utilised a broader band (100–500Hz), which may have been more effective than the narrow ones used by the BO in this study. Additionally,brain stimulation effects have been shown to be dependent upon numerous factors such as gender (Chaieb, Antal, and Paulus 2008; Russell et al. 2017), age (Bartrés-Faz, Rossi, and Pascual-Leone 2016) and brain-state (Cohen Kadosh, Soskic, et al. 2010; Neuling, Rach, and Herrmann 2013), indicating that a 'one-size-fits-all' approach like the one adopted in this chapter and the vast majority of the studies in the literature is not appropriate. In the case of chapter 3, participants who showed a positive effect at the optimal stimulation parameter, and therefore influenced the model to predict a beneficial effect of those parameters, may have benefited from the tRNS due to a combination of these personal variables. Whilst a recent study shed light on the combination of personal variables which might interact with the tRNS parameters (Sheffield et al. 2020), future work aiming to utilise BO to select stimulation parameters could attempt to include them in the protocol.

One clear way of individualising stimulation is to run the algorithm entirely on one person, thus eliminating any intra-individual variability (Lorenz et al., 2019). However, a limitation of this approach is that the optimisation must then be run on each individual that will undergo stimulation. One alternative approach, would be to include personal variables as a dimension of the BO in order to allow for

personalised Bayesian optimisation (pBO). As stimulation has been shown to have a differential effect according to participant-specific attributes, participants' personal variables could be inputted into the algorithm, allowing the pBO algorithm to identify the stimulation parameters most appropriate for an individual based on their personalised variables. This approach would enable personalised parameters to be selected for an individual, without having to re-run the BO, and could be extended to include further personal variables.

The variation inherent in human cognition can also be a problem for the BO algorithm. Sources of variation which are unaccounted for include gender (Chaieb, Antal, and Paulus 2008; Fumagalli et al. 2010; Russell et al. 2017), age (Bartrés-Faz, Rossi, and Pascual-Leone 2016) and genetics (Antal, Chaieb, et al. 2014; Teo et al. 2014). As previously mentioned, BO is usually applied in situations in which there are low levels of noise, a large amount of data, or both (Joy et al. 2016). Its application to human cognition exposes it to a situation in which there is a large amount of noise, and the amount of data is constricted due to the difficulties of testing large numbers of people. In the chapter 3, this combination of noise and limited sample size may have meant that the BO algorithm was unable to reliably identify a peak in the parameter space and instead modelled noise in the data. Previous studies in low-noise situations have suggested a sample size of 10D or 20D, where D is the number of dimensions present in the algorithm (Snoek, Larochelle, and Adams 2012; Hernández-Lobato, Hoffman, and Ghahramani 2014; Nguyen et al. 2017). More recent work employing BO in a noisy environment utilised 25D (Letham et al. 2019). I expected further variability and so collected 40D. However, this may still not have been enough. Tests of the equality of variance between the active and sham groups in the current chapter, revealed a trending difference

between the groups, suggesting that the distribution of participant performance may change according to stimulation parameter. This is an important factor to consider when designing the BO algorithm as the one employed in the current study assumes that variance is the same regardless of stimulation parameter. Allowing the variance to change with stimulation parameter would allow for a more effective algorithm. Taken together, these results suggest that attempts to minimise noise, as well as adapting the BO algorithm to cope with high levels of noise are vital.

Previous work investigating the neural correlates of this task have demonstrated the importance of the DLPFC in RITL (Cole, Bagic, et al. 2010; Cole, Etzel, et al. 2011; Cole, Braver, and Meiran 2017). In the current study DLPFC theta power and theta connectivity were investigated for their role in RITL. However, changes in DLPFC theta and DLPFC connectivity were not found in the present task when investigating trials with different levels of transfer. This suggests another frequency band in the DLPFC may play a more prevalent role in RITL. Previous work has also shown the role of the frontoparietal network (FPN) in RITL (Cole, Reynolds, et al. 2013), therefore parietal theta power was also investigated and effects of transfer distance were found during both the stimuli presentation and response periods in RITL. The analysis performed here showed decreased parietal theta power in trials with higher levels of transfer compared to trials in which the rules had been previously learnt, suggesting an increase in parietal theta with familiarity. This is in line with previous research which showed that parietal theta increases with how well an item is remembered (Jacobs et al. 2006).

In summary, in this study I validated the use of BO to determine parameters of tES in a complex cognitive task. In identifying the optimal stimulation parameters to improve transfer in RITL in Chapter 3, I initially found promising results as the

algorithm identified an optimal stimulation current intensity and frequency. However, the results of this chapter showed no effect of stimulation on transfer when utilising a large sample comparing active and sham stimulation. These conflicting results are likely due to the high levels of noise present in such a complex behaviour affecting the efficiency of the BO algorithm. Future studies should consider shifting from the inappropriate use of the one-size-fits-all approach to a personalised approach in parameter selection, and how noise can be further minimised. Such progress will eventually allow for more replicable stimulation, and successful translation to clinical settings. This approach will be discussed in Chapters 5 and 6.

Chapter 5

Personalising stimulation parameters using Bayesian optimisation for sustained attention

5.1 Introduction

The results of chapters 3 and 4 demonstrate that a one-size-fits-all approach to stimulation parameter selection is unlikely to work when manipulating complex human behaviours. The large amount of variability both in response to stimulation, and inherent in cognition, poses a significant challenge when utilising BO for this purpose. However, a number of studies have sought to adapt stimulation parameters on an individual basis to maximally improve the performance of the individual. One clear example of this is the use of an individual's peak alpha frequency, i.e. the frequency within the alpha range at which there is the highest power, to guide frequency selection in tACS (Clancy et al. 2018; Fresnoza et al. 2018; Vosskuhl, Huster, and Herrmann 2016). The idea of personalising stimulation parameters can also be applicable to the use of BO in parameter selection. Previous work has

aimed to derive personalised tACS parameters to induce phosphenes by running all of the iterations of the BO algorithm within one individual (Lorenz, Simmons, et al. 2019). Whilst this approach does achieve entirely personalised parameters and is ideal for simple objective functions, it is impractical when applying stimulation and BO to more complex behaviours in which there are larger levels of noise. Specifically, increased variability in behavioural responses requires a paradigm in which a larger amount of data is collected to achieve a reliable behavioural metric. Such a paradigm, when used to test a large number of different stimulation parameter combinations, would take far too long and leave the participant exhausted and unable to focus on the task, therefore not giving usable data. Additionally, when aiming to utilise tES in the clinical setting, patients are often unable to receive repeated stimulation, making a within-subject design unusable. An alternative way to personalise stimulation parameters using BO, would be to incorporate personal variables such as IQ or age into the algorithm, therefore allowing the algorithm to suggest parameters based on specific values of the personal variables. This could be done by extending the GP to incorporate the personal variables, and so treat them as an extra dimension in the search space, with the constraint that the BO cannot control the value of the personalised variable to be measured as it is specific to the individual being tested (Figure 5.1). When applied to human cognition, the personal variables could be any measurable attribute that varies amongst individuals, but would ideally characterise the highest amount of variability in response to stimulation as is possible. Examples of such a personalised variable could include gender, IQ, ability in a skill of interest such as working memory capacity, and neural correlates of the specific cognitive function such as EEG frequency band measurements or fMRI BOLD signals.

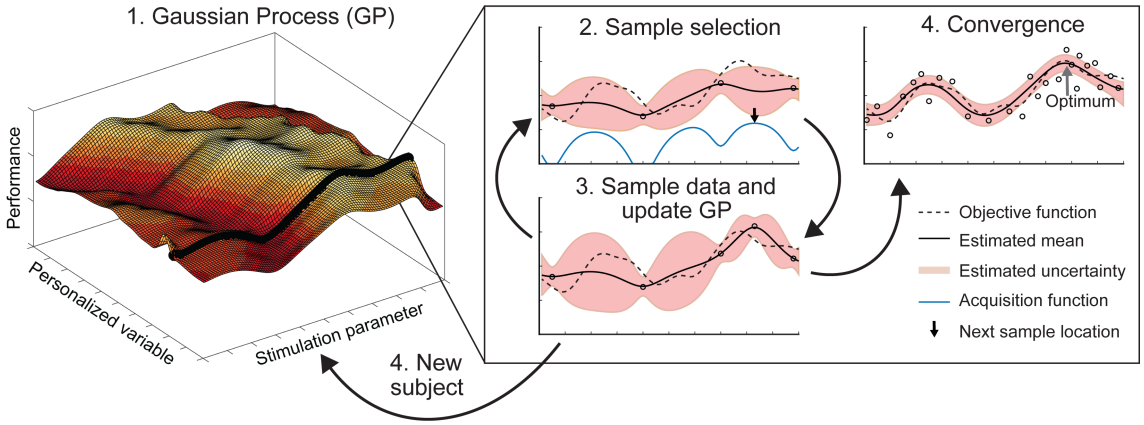


Figure 5.1: **An overview of the pBO process** 1.The Gaussian process models the existing data, including a personalised variable which is measured for the individual. 2. An acquisition function is fitted for the specific value of the personalised variable and selects where to sample next. 3. new data is sampled. 4. The GP is updated or convergence occurs according to pre-defined criteria.

Using such a method, an initial study could be run, collecting data from individuals spanning the range of the personal variables selected, during which the algorithm would identify 'optimal' stimulation parameters at each value of the personal variable. After this initial study had been run, any new individuals requiring tES would simply need their personal variables measured, and the algorithm could recommend stimulation parameters tailored to that individual. This personalised Bayesian optimisation (pBO), although requiring a substantial initial study, would avoid the need to re-run the BO algorithm in each new individual to be tested, making it ideal for commonly used behavioural paradigms as well as in a clinical setting. A similar BO method has previously been utilised in the machine learning literature (A. Krause and Ong 2011), in this case called contextual BO, and has been applied to areas such as fraud regulation (Hu et al. 2019), the control of wind farms (Park 2020) and cancer treatment discovery in mouse models (Durand et al. 2018). One recent example utilised a contextual GP to model partic-

ipant's behaviour in a learning and decision-making task (Schulz, Konstantinidis, and Speekenbrink 2018).

When applying pBO to tES parameter selection, it is vital to select a personalised variable that will account for a large portion of the variability seen in response to stimulation. One promising behaviour in this respect is sustained attention, as previous work has identified the EEG marker theta:beta ratio (TBR) as being a likely predictor for an individual's response to tRNS (Harty and Cohen Kadosh 2019). As the name suggests, TBR is the ratio of theta to beta power in an individual's EEG, measured in this instance during a resting state recording. Previous work has shown the involvement of TBR in sustained attention by demonstrating that it is elevated in clinical groups with sustained attention deficits such as ADHD (Arns, Conners, and Kraemer 2013). Interestingly, the same study also demonstrated that increased levels of TBR are associated with a more positive response to interventions aimed at improving sustained attention such as neurofeedback paradigms (Arns, Conners, and Kraemer 2013). More recent work has investigated the effects of tRNS on sustained attention by delivering 1mA, 2mA or sham tRNS to the fronto-parietal network during a sustained attention task. It was found that 1mA tRNS was associated with improvements in sustained attention compared to sham, whilst 2mA was not, in line with stochastic resonance mechanistic theories of tRNS. Crucially, participants' baseline TBR moderated the effects of 1mA tRNS, with individuals with higher TBR responding more positively to tRNS (Harty and Cohen Kadosh 2019). It should be noted that this study used a broad frequency band (100-500Hz), in comparison to the frequency bands used in the current chapter. Building on this previous work which suggests that TBR would be a suitable candidate for a personalised variable in the pBO algorithm, in this

chapter, I aim to implement pBO in order to identify the 'optimal' tRNS parameters for improving sustained attention whilst incorporating TBR as a personalised variable.

5.2 Methods

5.2.1 Participants

Thirty-eight healthy right-handed participants (22 females) between the ages of 18 and 47 years (mean± sd = 25.10±6.35) completed the experiment. Data of one participant was rejected for low accuracy (responding to under 50% of trials in 3 out of 4 blocks). In addition, two participants' second session data was not used as one failed to attend the session and another participant's baseline data was corrupted and so could not be used. All participants gave informed consent and were compensated for their time.

5.2.2 Procedure

The setup used in this experiment was consistent with Harty and Cohen Kadosh 2019 with the addition of the BO algorithm. Participants completed two sessions of testing on consecutive days. Each session started with one resting-state EEG recording, which, on the first session was followed by 14 practice trials of the task. Participants then completed 4 baseline blocks of a sustained attention task in which stimulation was not administered, followed by 4 blocks in which stimulation was administered. Prior to each stimulation block, the pBO algorithm was run to determine the tRNS parameters which were then delivered to the participant whilst

they completed the task. After each stimulation block, participants answered a questionnaire aimed at determining the degree of sensations experienced during stimulation (Figure 5.2).

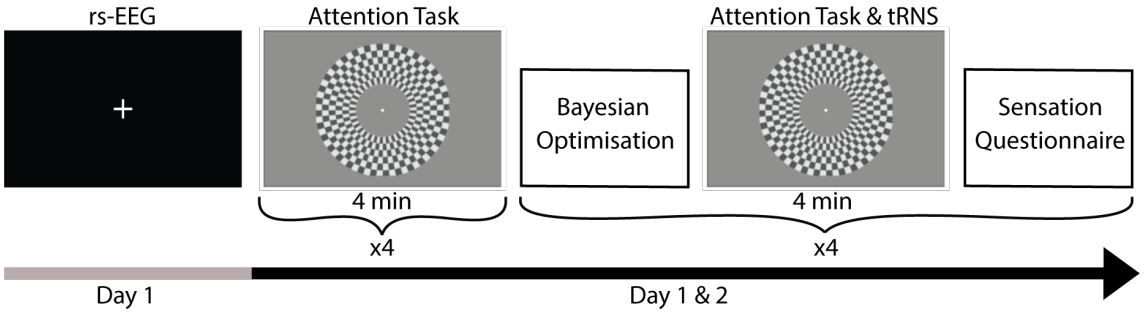


Figure 5.2: **An illustration of the experimental paradigm** Participants underwent a rs-EEG recording and 4 blocks of the sustained attention task. The pBO algorithm was then updated with the participant's personalised variables and another block of the task was run, after which a sensation questionnaire was completed. The pBO algorithm and the task were performed 4 times.

5.2.3 Behavioural task

The sustained attention task consisted of a fixation cross placed in the middle of the screen which participants were instructed to fixate on. A stimulus consisting of a hatched circle was placed around the fixation cross. The stimulus' baseline contrast level was set at 65% and at each target it would decrease to 35% contrast over 1.6 seconds (target), before returning to baseline contrast over 0.4 seconds. The time between targets was randomly selected to be 4, 7.2 or 10.4 seconds. Participants were instructed to respond to any gradual decreases in contrast of the stimuli as quickly as possible using a button press. Each of the eight blocks completed by participants lasted four minutes and included 25 targets. Participants had a break of approximately 2-4 minutes between blocks. A task block was excluded if

participants failed to respond to at least 50% of the targets.

5.2.4 Behavioural performance

Performance was evaluated by quantifying participants' variation in response times over each task block. To do this, a modified version of the Coefficient of Variation (CoV) calculation previously used by Fortenbaugh et al. (2018) and Harty and Cohen Kadosh (2019) was used. The modified version was used as the objective function for the pBO algorithm and measured the standardised CoV improvement from baseline to stimulation block:

$$CoVI = \frac{CoV(block_s) - mean(CoVblock_{bl})}{sd(CoVblock_{bl})} \tag{5.1}$$

where

$block_s$ is the stimulation block of interest

$block_{bl}$ are the baseline blocks 1-4

5.2.5 Stimulation

tRNS was delivered using the Novostim brain stimulator (Innosphere Ltd, Berlin). Square brush electrodes (area 25cm^2) were soaked in 0.45% saline solution and were positioned at locations F4 and P4 according to the 10-20 system and held in place using a neoprene cap. Prior to the first task block including stimulation, a short resistance test of 30 seconds ramp-up and 30 seconds ramp-down was conducted to ensure resistance was between 0.5kΩ and 1kΩ. tRNS current intensity and frequency were chosen by the pBO algorithm. Current intensity ranged from 0mA to

3mA peak-to-peak in steps of 0.1mA and frequency was arranged in bins of 100Hz from 100-200Hz up to 900-1000Hz with a 50Hz overlap (i.e. 100-200, 150-250, ... 850-950, 900-1000Hz). As no effect of stimulation was expected at lower currents, but in order to include sham at 0mA, the algorithm was stopped from sampling between 0.1mA and 0.5mA. During active stimulation participants received a 30-second ramp-up period, 4 minutes of stimulation followed by a 30-second ramp-down period. Sham stimulation (0mA) consisted of a 30-second ramp-up period to 1.5 mA peak-to-peak with a random frequency, followed by an immediate 30-second ramp-down.

5.2.6 Theta/Beta Ratio (TBR)

In order to calculate TBR during the experimental session and then incorporate it in the pBO algorithm, immediately after acquiring the resting state EEG recordings, the raw EEG data was automatically pre-processed using custom made scripts in the EEGLab toolbox (Delorme and Makeig 2004) in Matlab (version R2018b). A linear 40Hz low-pass filter was applied, after which data from channel Cz was epoched into two-second segments. Noisy epochs with electrical potentials above 100μV were rejected and a Fast Fourier transformation (FFT) was performed. The absolute theta (4-7.5Hz) and beta power (14-30Hz) were normalised to the absolute power within the 1.5-30Hz band. The TBR was then calculated by dividing the normalised theta power by the normalised beta power. TBR was then inputted into the BO algorithm for each participant as a personalised variable.

5.2.7 Personalised Bayesian optimisation (pBO)

The pBO algorithm was retrained prior to each stimulation block using all previously gathered data, as well as the TBR value for the upcoming block. The algorithm then outputted the stimulation parameters for this upcoming stimulation block. Inputs into the algorithm included tRNS current intensity and frequency, with participant TBR inputted as a personalised variable. The algorithm aimed to identify the stimulation parameters most capable of improving sustained attention based on each participant's TBR.

The initial burn-in phase of the pBO algorithm consisted of data from 18 participants, in which the stimulation parameters were randomly sampled, resulting in 124 behavioural data points. After the burn-in, the pBO algorithm was initialised in order to select stimulation current intensity and frequency parameters, whilst incorporating TBR measures. The pBO algorithm was run on a further 20 participants, resulting in 273 behavioural data points in total, 124 from the burn-in and 149 from the optimisation phase.

The GP used in the pBO algorithm had a constant basis function and a Mátern 3/2 automatic relevance determination (ARD) kernel with the property to adjust the fit to the spread of the data. The pBO algorithm utilised a UCB acquisition function with importance sampling to guide the parameter selection process. The output of the pBO algorithm was encoded and automatically loaded into the stimulation software to maintain double blinding. Additionally, the pBO algorithm was coded in such a way that if there was no clear pattern in the data, and therefore the acquisition function was unable to select a parameter combination to sample, as all points were equally rated, the algorithm would revert back to random sampling.

5.2.8 Sensation Questionnaires

A questionnaire was adapted from Fertonani, Rosini, et al. 2010 to enable participants to rate the sensations induced by tRNS on a 10-point Likert scale. Sensations rated included: "itching", "pain", "burning", "warmth/heat", "fatigue/decreased alertness", "flashing lights (other than the stimulus)" (see Appendix A.2 for questionnaire).

5.3 Results

5.3.1 Performance of the pBO algorithm

Whilst selecting stimulation parameters to test, the pBO algorithm reverted to randomly selecting parameter combinations, as there was no clear pattern in the data. Data collection was halted after 38 participants were tested, with at least the last 5 participants having been given random stimulation parameters. Figure 5.3 shows the current x frequency experimental space, regardless of TBR. As can be seen, the surface was relatively flat and no clear peaks were found.

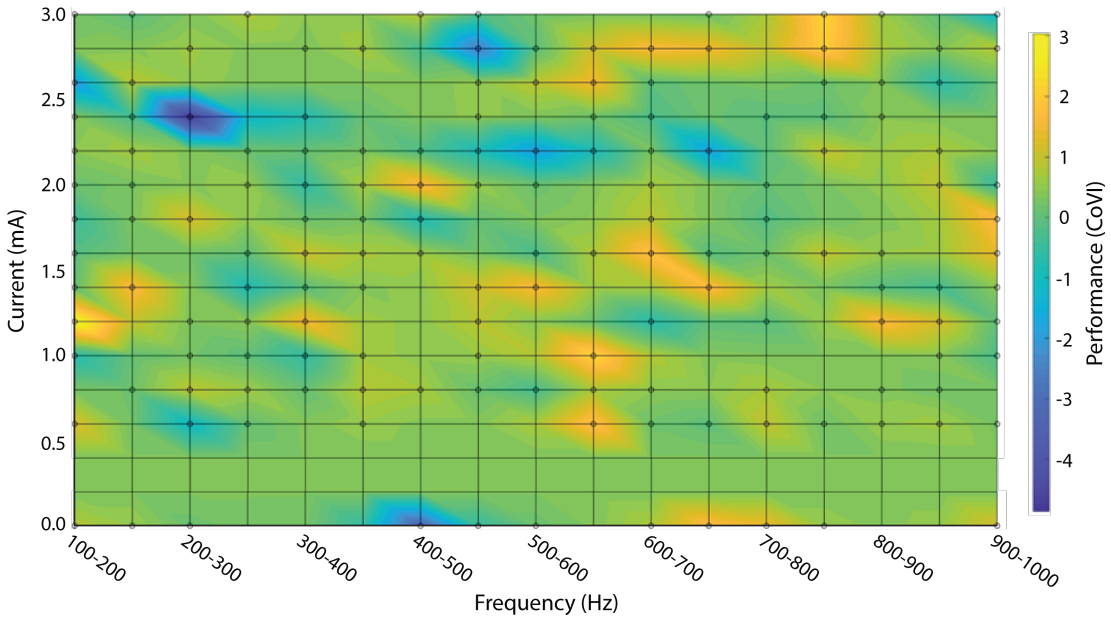


Figure 5.3: **Performance in the sustained attention task** The estimates of performance in response to different tRNS current intensities and frequency bands, regardless of TBR

To investigate each of the pBO dimensions individually, each participant's COVI was plotted against the stimulation parameters used in that block. As Figure 5.4 shows, there was no effect of current, frequency or TBR on COVI. Figure 5.4c also demonstrates that the TBR values were heavily skewed towards the lower end of the range.

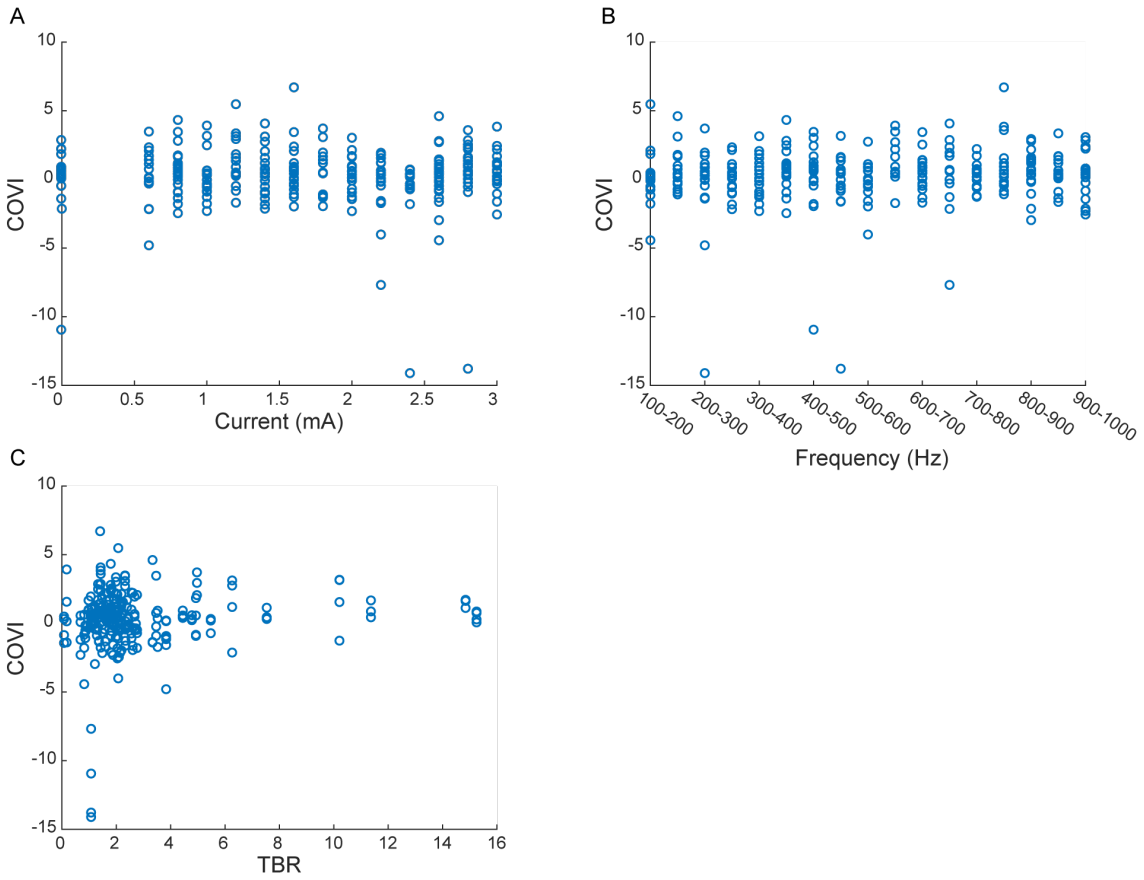


Figure 5.4: **Performance in the sustained attention task according to each input variable.** The scatter plots demonstrate that there is no change in COVI according to A) Current, B) Frequency or C) TBR.

5.3.2 Sensation analysis

In order to investigate the effects of varying tRNS current and frequency on participant reports of sensations during stimulation, linear mixed effect models were run with current and frequency as fixed factors and participant, session and block as nested random factors. This analysis showed an increase in participants' reports of itching, pain, burning and warmth in response to higher currents, as well as a decrease in burning sensations with increased stimulation frequency. Full

results of the analysis, including p-values corrected for multiple comparisons can be found in Table 5.1.

Variable	Predictors	Estimates	CI (95%)	df	t-value	p-value
Itching	Current	0.315	0.207–0.422	192	5.77	**<.001**
	Frequency	-0.092	-0.201– -0.092	192	-1.68	.113
Pain	Current	0.162	0.051–0.272	192	2.88	**.009**
	Frequency	-0.101	-0.213–0.011	192	-1.79	.097
Burning	Current	0.155	0.040–0.271	192	2.65	**.016**
	Frequency	-0.208	-0.324– -0.092	192	-3.53	**<.001**
Warmth	Current	0.139	0.034–0.243	192	2.61	**.016**
	Frequency	-0.102	-0.207–0.004	192	-1.89	.090
Fatigue	Current	0.002	-0.096–0.100	192	0.04	.967
	Frequency	0.014	-0.086–0.113	192	0.27	.832
Lights	Current	0.100	-0.009–0.211	192	1.81	.097
	Frequency	0.066	-0.044–0.176	192	1.18	.270

Table 5.1: Results of the mixed effect models investigating the effects of tRNS current and frequency on sensations experienced by the participant. p-values were computed using Benjamini-Hochberg corrections for multiple comparisons

5.4 Discussion

In this chapter I aimed to identify the optimal tRNS parameters for improving sustained attention whilst accounting for an individual's baseline TBR. In doing this, I hoped to be able to personalise the stimulation parameters participants receive simply through measuring their TBR. However, after 38 participants had been tested, and 273 stimulation blocks completed and used in the training of the pBO algorithm, the algorithm had reverted to randomly sampling tRNS parameter combinations as there was no clear pattern in the data. These results could be explained very simply by there being no effect of tRNS on sustained attention. However previous work by Harty and Cohen Kadosh (2019) using a well powered

study, suggested that there was an effect at 1mA. It could be that there is a small effect and given a larger sample, it would have been detected. This may be the case, especially as the pBO algorithm detects peaks in the experimental space, which is easier if they are broader peaks, with higher amplitudes. As Harty and Cohen Kadosh (2019) showed an effect of 1mA but no effect of 2mA, it may be that there is only a narrow peak to detect in this case, making it harder for the pBO algorithm. Additionally, as shown by the simulation analysis in chapter 3, BO algorithms are not as effective in environments with large levels of noise. As can be seen in Figure 5.4, there is a considerable level of noise in the measures used in this experiment, potentially limiting the ability of the algorithm to identify peaks. The use of a broad frequency band in Harty and Cohen Kadosh (2019), in comparison to the narrow ones used in the current chapter could also contribute to the conflicting results between Harty and Cohen Kadosh (2019) and the current work. It may be that the broader band is needed to induce an effect, leading to no effect in the narrower band, or that the narrower band missed the crucial frequencies responsible for the behavioural effect. An alternative approach when using BO that could account for this, would be to allow the BO algorithm to select where the frequency band is centred and how wide it is. This would make the algorithm more complicated, but would identify whether broad or narrow frequency bands are most suitable.

Another potential explanation for the present results is the use of TBR as a personalised variable was unfortunately unsuccessful. Previous research had suggested that TBR was important in sustained attention (Arden et al. 2010) and would moderate the effects of tRNS (Harty and Cohen Kadosh 2019). However as there were no effects to moderate, this was not the case. It may be that the effects of tRNS on sustained attention are only detectable in samples with higher power and cannot

be detected by optimisation algorithms relying on finding trends in the data. Additionally, it may be that the involvement of TBR in sustained attention does not account for enough variability in the sustained attention behavioural measures so as to impact the tRNS parameters required to modulate performance. Whilst TBR is one of the most studied EEG correlates, utilising a cognitive marker such as this may be adding unnecessary complexity to an optimisation algorithm which already has large levels of noise to deal with to identify an optimum. Therefore, it may be that the use of alternative personalised variables is desirable. Ideally these variables would be as simple as possible, whilst accounting for as much of the behavioural variability as possible. For example, a measure of participant's baseline ability in the cognitive task of interest or of the cognitive domain in general may be a viable option. Alternatively, variables such as gender and age could also be used as they require minimal effort to measure, whilst potentially accounting for additional variability. Analysis of the sensation data collected in this experiment revealed that increased tRNS current intensity increased participants' ratings of a number of sensations, whilst increased tRNS frequency decreased reports of burning sensations. Whilst tRNS has previously been shown to be well tolerated (Schecklmann et al. 2017), these results show that increasing the current intensity has the potential to undermine the blinding of stimulation conditions, in agreement with previous work showing the potential ineffectiveness of blinding in tES protocols (Bereznicki et al. 2019; Fonteneau et al. 2019; Greinacher et al. 2019). Blinding was not examined directly in this study as only a small number of participants received sham tRNS at 0mA. Had the pBO algorithm identified a peak that was not the sham stimulation at 0mA, then it would have suggested that the selected point would be better than sham stimulation as the pBO algorithm had not converged at 0mA.

Whilst the pBO algorithm is theoretically promising, in this study a number of decisions were made that may have contributed to the failure of the algorithm to identify optimal tRNS parameters. Key amongst these is the use of TBR as a personalised variable. In the following two chapters I will aim to expand the use of pBO in parameter selection, whilst using a more simplistic personalised variable.

Chapter 6

Personalising stimulation parameters using Bayesian optimisation for arithmetic skills

6.1 Introduction

This chapter builds on the work from chapter 5 by utilising a pBO algorithm to provide predictions of the best stimulation parameters for an individual based on a personalised variable. I utilise an arithmetic behavioural paradigm due to its simplicity and the ability to complete a lot of trials in a short space of time, allowing for more reliable measures. Arithmetic processes have also been shown to be highly important in the success of one's future career and socio-economic status (Ritchie and Bates 2013) and arithmetic impairment has been associated with acquired and congenital brain disorders (Willmes 2008; Cohen Kadosh and Walsh 2007). Previous work has shown a wide variety of skills are needed for solving arithmetic problems in both typical and atypical populations (Kaufmann et al. 2013; Parsons and Bynner 2005). Interestingly, a recent study on arithmetic skills utilised fNIRS to show that in individuals with higher mathematical abilities there was increased

activation in frontal brain areas whilst performing arithmetic operations, whereas there was no such increase in low performers (Artemenko et al. 2015). This finding suggests that a suitable personalised variable for the pBO algorithm may be a measure of participant's baseline arithmetic ability as it would indirectly account for these different cognitive processes happening in high/low performers, without the need for a complex and lengthy neuroimaging setup and recording. Further work has implicated the left frontalparietal network in arithmetic processing (Ischebeck et al. 2007; Klein et al. 2016; Matejko and Ansari 2015; Nieder 2016) and has shown that most arithmetic-related frequency band changes occur in the theta and alpha bands (Grabner and De Smedt 2012; Hinault and Lemaire 2016). These findings suggest that tACS would be the most suitable type of tES to use here as these frequencies are easily targeted using it.

Building on this previous work, I aim to identify the optimal tACS frequency and current intensity for individuals across the range of a personalised variable. In the previous chapter I utilised TBR as the personalised variable, however this proved to be ineffective. In this chapter I utilise a more simplistic personalised variable in the form of baseline performance in the arithmetic task. The use of baseline performance as a personalised variable allows for faster and easier acquisition of the personalised variable compared to a potentially lengthy EEG setup when measuring TBR in chapter 5. Additionally, baseline measures are more directly related to the measures of participant task performance. Therefore, it is hoped that a larger degree of intra-individual variability will be accounted for, and therefore more accurate parameters will be generated.

6.2 Methods

Fifty right-handed participants aged 18-30 years old (mean±std age = 22.52±4.09, 31 females) took part in this experiment. All participants gave informed consent and were compensated for their time. During the experiment, participants' resting-state EEG (rs-EEG) was recorded, and participants then performed ten practice trials of the arithmetic task. Participants then completed four blocks of 50 multiplication problems, the first of which was used to assess baseline arithmetic abilities, and the other three blocks were completed whilst receiving tACS. Prior to each tACS block, the BO algorithm was run to determine the stimulation parameters (current and frequency) to be delivered during the upcoming experimental block based on the participant's baseline performance. rs-EEG was recorded after each stimulation block (Figure 6.1b).

6.2.1 Behavioural Task

The behavioural task was comprised of blocks of 50 novel multiplication trials, with each block divided into 5 mini-blocks of 10 trials each, with a 30 second break between each mini-block. Multiplications consisted of a single-digit number multiplied by a two-digit number, with a three-digit answer. To limit variations in difficulty, none of the multiplication problems included operands with the digits 0, 1, or 2 . In addition, the two-digit operand was not smaller than 15, did not use repeated digits, and was not a multiple of 10. The problem was presented on screen with a correct and incorrect answer positioned underneath them on the right and left sides of the screen. The position of these correct and incorrect answers was randomised and the digits always differed by ten. There was no limit for how long

a participant could take to respond, however they were instructed to respond as quickly and as accurately as possible. Participants responded to a trial by pressing a button on the left or right of a response pad which corresponded to the left and right answers on the screen (Figure 6.1a). Each problem was presented only once, and their order was randomized. Multiplication difficulty was quantified by the magnitude of the answer, with difficulty being equally distributed across blocks.

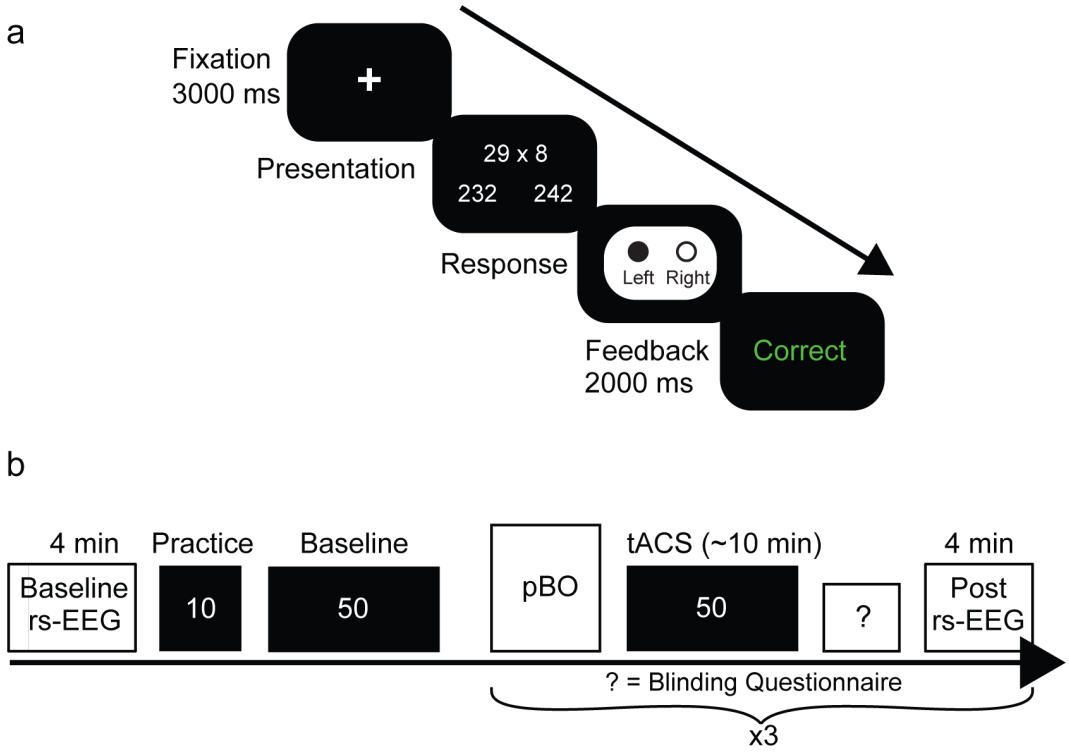


Figure 6.1: **An illustration of the arithmetic paradigm. a)** shows the arithmetic task. **b)** shows the components of each session.

6.2.2 Measurement of Behavioural Performance

Participants' arithmetic abilities were quantified in terms of response times and accuracy through the use of drift rates computed by two-choice EZ-diffusion mod-

elling. This allowed for the component of interest, the drift rate, which reflects ability, to be targeted. The EZ-diffusion model was chosen as it allows a relatively simple way to model the response time and accuracy to be combined in to one outcome that could be optimized through the pBO procedure. Drift rates in the stimulation blocks were normalised to the baseline block. To eliminate any dependency between a participant's baseline ability and the normalised score, baseline trials were randomly divided into half and two separate drift rates were calculated. One of these was used as a participant's baseline ability score, whilst the other was used to normalise the participant's performance in the stimulation blocks.

6.2.3 Experimental Procedure

Participants' baseline rs-EEG were initially recorded, after which the task was explained to them and they had an opportunity to ask questions. Participants then completed 10 practice trials of the arithmetic task and any further questions this raised were answered. Participants then completed a baseline block of 50 trials, with the resulting drift rates being used to normalise the stimulation blocks and provide a baseline drift rate score for the personalised variable. The pBO algorithm was then run with this baseline drift rate as a variable. The algorithm determined which tACS frequency and current intensity the participant would receive, and this was automatically initiated when they started the next block of 50 trials. In this manner a blinded protocol was maintained. Once the task block was complete, participants answered a blinding/sensation questionnaire and then a post rs-EEG recording was performed. In total participants completed three stimulation blocks, with the pBO algorithm being run individually for each one (Figure 6.1b).

6.2.4 pBO Algorithm

Prior to each stimulation block, the BO algorithm was run. Overall, this consisted of an initial burn-in phase in which the first 60 stimulation blocks across 20 participants were assigned random stimulation parameters, enabling the GP to be trained on some data. After this, the pBO phase was initiated, with the pBO algorithm being retrained prior to each stimulation block in order to determine the stimulation parameters to be used in the upcoming stimulation task block. Stimulation current intensity and frequency were used as inputs in the pBO algorithm. Additionally, participant baseline ability was entered as a personalised variable and the algorithm aimed to determine the stimulation current and frequency which elicited the best performance in arithmetic abilities given a participant's specific baseline ability. It should be noted that the algorithm was trained on all data previously collected across participants, and each participant did not have their own pBO algorithm. A stopping criteria of 50 participants was set in order to halt the algorithm.

6.2.5 Transcranial Alternating Current Stimulation

The alternating current stimulation was administered over the left frontoparietal network. The tACS was delivered via two stimulation (3.14 mm diameter) NG Pistim Ag/AgCl electrodes (F3 and P3) with one return electrode (Cz) using the Starstim 32 (Neuroelectrics, Barcelona). The impedances of the electrodes were held at <10 kΩ. The stimulation intensity ranged between 0.1-1.6 mA peak-to-peak in steps of 0.1 for the burn-in phase of the study. For the optimization, 0 mA was added to control for possible sham influences. The maximum stimulation

intensity was chosen based on a small pilot study on three subjects to determine the maximum comfortable intensity. The stimulation frequency ranged between 5-50 Hz in steps of 1 Hz for the whole experiment. Stimulation was administered in a double-blind manner during the three experimental blocks with a maximum of 10 minutes for each block. Stimulation started 45 s before the start of the block and changed after every block. If the subjects received a stimulation intensity of 0 mA during a block (sham stimulation), a ramp-up and a ramp-down of 30 s was initiated to provide the initial skin sensations during stimulation to ensure blinding. When the subject completed a block within 10 minutes, stimulation was ramped down for 30 s and the subject proceeded to the four-minute rs-EEG. Note that in cases where subjects completed the task in under 10 minutes, they did not receive the full length of stimulation. These stimulation blocks did not differ from the stimulation block in which the subject received the full length of stimulation except for performance. Twenty-four of the 150 stimulation blocks had a duration of less than 10 minutes but more than 7.50 minutes, and 126 stimulation blocks had a duration of more than 10 minutes. This posed no problem, since the present study only investigated the online effects of tACS.

6.2.6 Resting State-EEG Recordings and Pre-processing

Resting state-EEG recordings were made at the start of the study (before baseline measurements) and immediately after every stimulation block. Electrophysiological data were obtained with eight gel Ag/AgCl electrodes (F3, P3, F4, P4, Fz, Cz, Pz, AF8) according to the international 10/20 EEG system using the Starstim 32 sensor system (Neuroelectrics, Barcelona, Spain), with no online filters. The

ground consisted of adhesive active common mode sense (CMS) and passive driven right leg (DRL) electrodes which were positioned on the right mastoid. All EEG measurements had a duration of four minutes in which the subjects had their eyes open while watching a fixation point in the middle of the screen. Raw EEG data were recorded and used for offline analysis using EEGLAB (Delorme and Makeig 2004), in Matlab R2018b (Mathworks). Data were high-pass filtered at 0.1 Hz, and low-pass filtered at 50 Hz. Visual inspection was carried out to remove artefacts caused by muscle movement. A recording was rejected from the analysis if more than 25 percent of the data in a given block was removed. This resulted in four rejected datasets (2% of the data). Independent component analysis was used to remove blinks and noisy components. On average, 1.23 ± 0.62SD components were rejected per subject, with a maximum of three components and a minimum of zero.

6.2.7 Sensation Questionnaire

After completing a stimulation task block, the subjects filled out a questionnaire in which they were asked several questions designed to gauge the level of sensation experienced during stimulation. Sensations were rated on a scale of 1 to 10. Sensations included: "itching", "pain", "burning", "warmth/heat", "fatigue/decreased alertness", "flashing lights (other than the stimulus)" (see Appendix A.2 for questionnaire).

6.3 Results

6.3.1 Performance of the pBO algorithm

Overall performance of the pBO algorithm can be seen in Figure 6.2, in which the GP was retrained after each iteration, including those in the burn-in phase to estimate the best predicted performance regardless of baseline ability. As can be seen in the figure, the best predicted performance in response to stimulation increases slightly during the burn-in phase as better parameters are randomly chosen. During the pBO phase, the algorithm increasingly finds stimulation parameters with better responses, and also re-tests those parameters, signified by the decrease in the predicted performance following an increase.

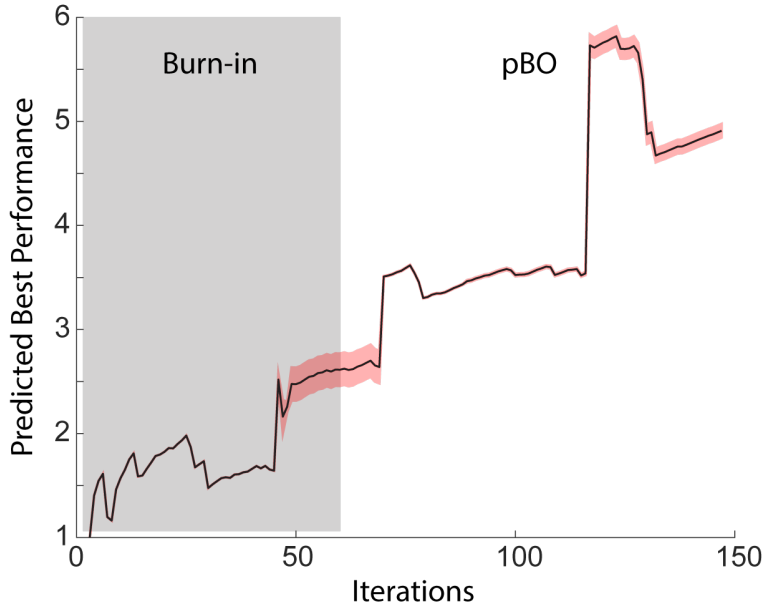


Figure 6.2: **Best predicted performance of the pBO model at each iteration.** shaded areas denote the uncertainty at the point of best predicted performance.

6.3.2 Behavioural results

Predicted performance by the GP, in response to different stimulation parameter combinations, can be seen in Figure 6.3, at different baseline abilities. As can be clearly seen, the stimulation parameters that are predicted to have the best behavioural response vary with baseline ability (mean-1SD (38.67 Hz, 0.97 mA), mean (16.67 Hz, 0.88 mA), mean+1SD (18 Hz, 0.6 mA)), suggesting an effective use of baseline ability as the personalised variable.

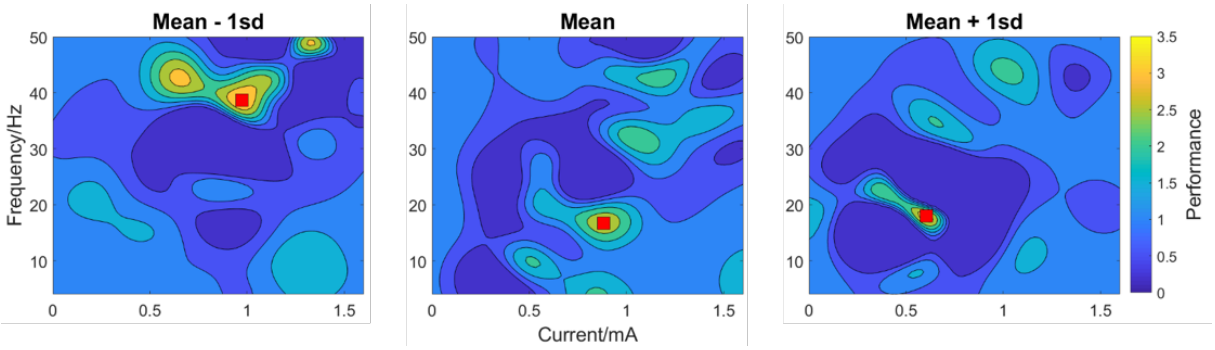


Figure 6.3: **Results of the pBO algorithm at mean-1sd, mean and mean+1sd of baseline ability.** Best stimulation parameter denoted by the red square.

A more detailed version of these results can be seen in Figure 6.4, where participants with lower baseline ability scores are predicted to respond best to higher current intensities and frequencies.

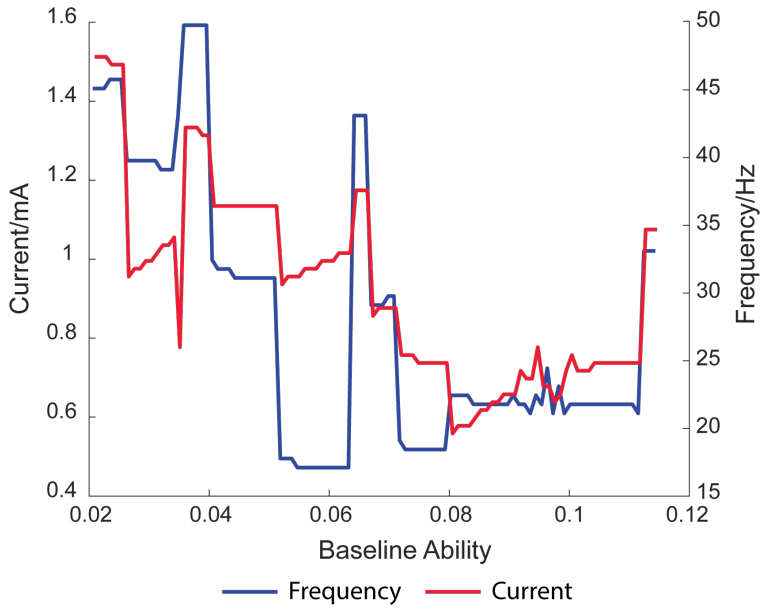


Figure 6.4: **Results of the pBO algorithm showing the best found current and frequency at each baseline ability**

There was no clear effect of baseline ability on the predicted amount that participants could improve due to the best stimulation parameters for their baseline ability (Figure 6.5a). However, participant baseline abilities were skewed towards the lower end of the range (Figure 6.5b), suggesting that predictions at higher baseline abilities are less reliable.

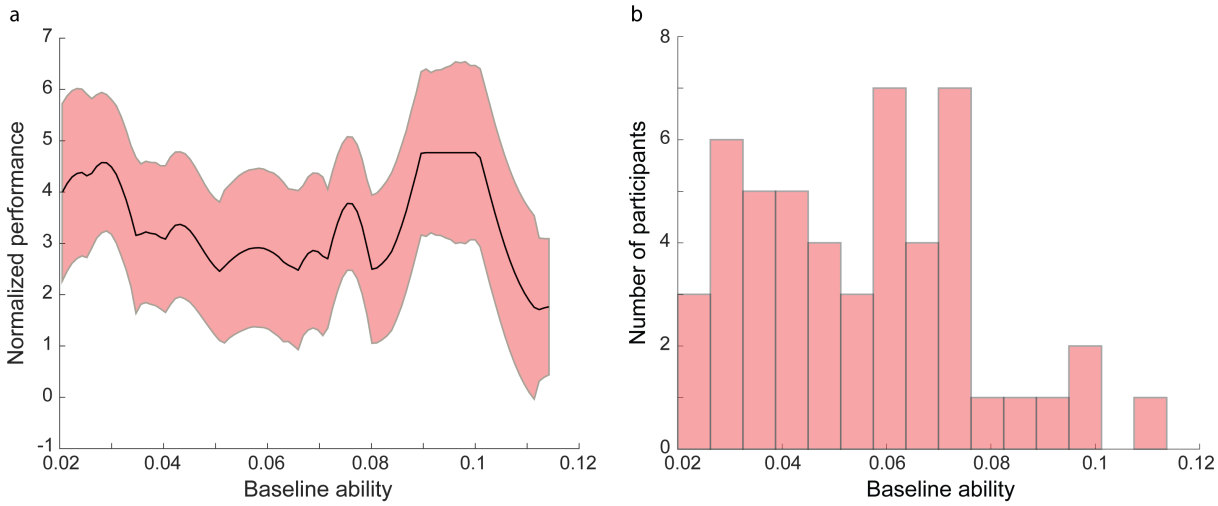


Figure 6.5: **Maximum predicted performance at different baseline abilities and the number of participants at each baseline ability. a)** the performance predicted by the GP at the best predicted parameter combinations for each baseline ability. **b)** the distribution of participants across the baseline ability range, note the increased number of participants at lower baseline abilities compared to higher baseline abilities.

6.3.3 Simulation analysis of pBO performance

Simulation analysis examined the effectiveness of the pBO algorithm at various levels of noise. For this a 3-dimensional Hartmann surface was used, a commonly used test surface in the machine learning literature, where each dimension was scaled between 0 and 1. The pBO algorithm was compared to random sampling of the surface, and a 'standard' BO algorithm. In order to accurately simulate these algorithms on the 3-dimensional test surface, both the pBO and BO algorithms optimised two of the dimensions (frequency and current), whilst the location to be sampled on the third dimension (personalised dimension) was randomly selected. The pBO algorithm suggested where to sample in the first and second dimensions conditioned upon the third dimension, whilst the BO algorithm suggested the first

and second dimensions independently of the third dimension (i.e. it was not given any information on the third dimension). In this way, the third dimension was used to influence the optimisation in the pBO algorithm, but not the BO. Each noise level was run 30 times for each algorithm type and the results averaged and are shown in Figure 6.6. Results showed that when there was no noise introduced to the system, the pBO algorithm clearly outperformed both the random sampling and BO algorithms, converging faster, and finding higher values (Figure 6.6a), as well as being more accurate and identifying optima closer to the true optima of the Hartmann surface (Figure 6.6b). As the noise levels increased, the difference between the pBO algorithm and the random sampling and BO gradually decreased. However random sampling and BO never outperformed the pBO algorithm.

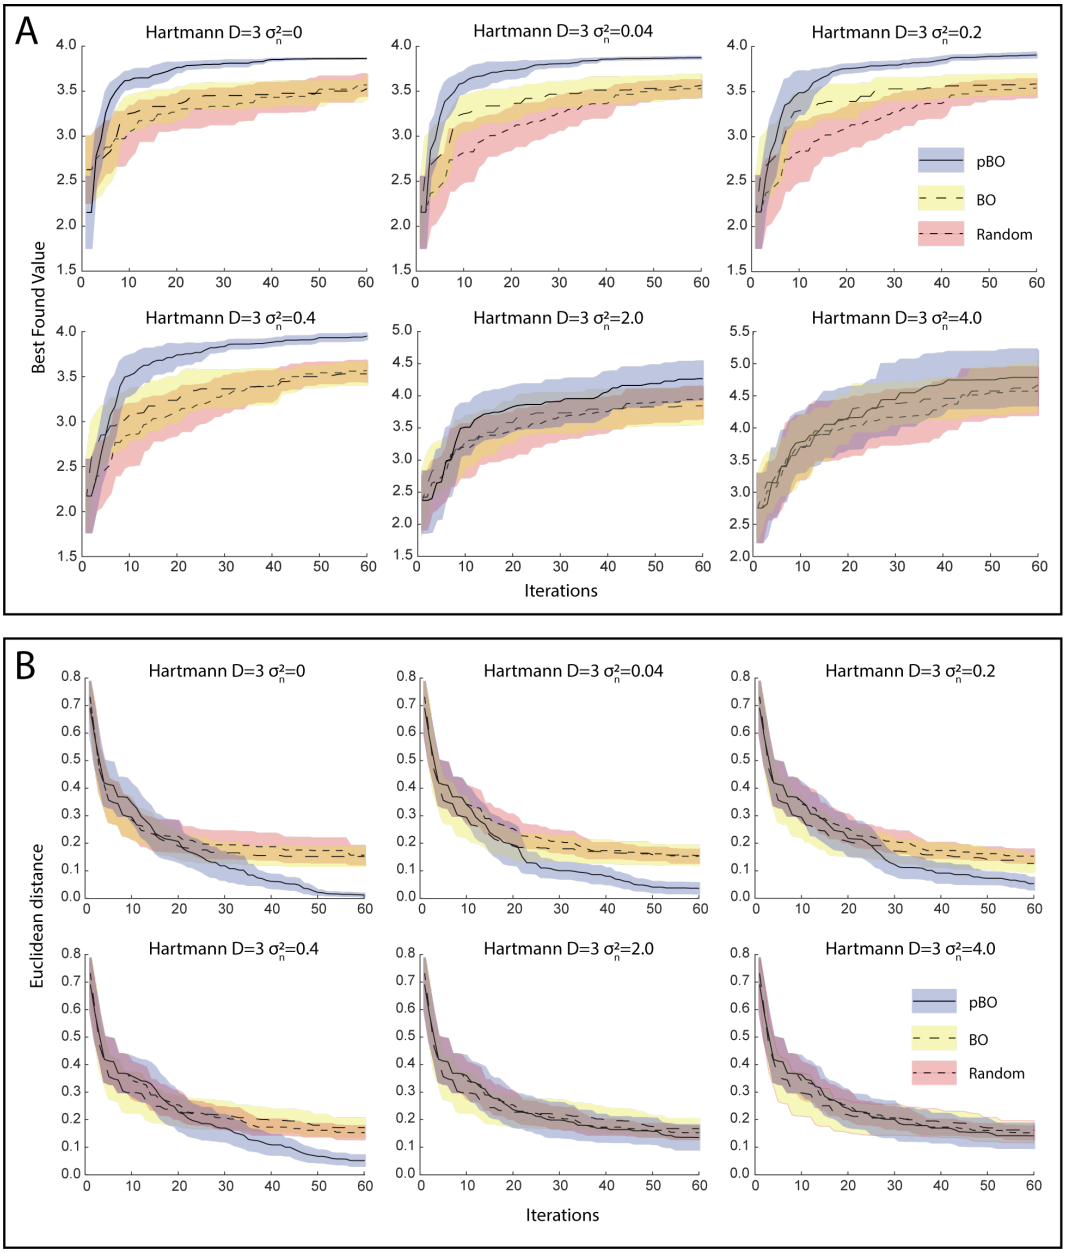


Figure 6.6: **Simulation results on the 3-dimensional Hartmann surface comparing pBO, BO and random sampling at various levels of noise (σ). A)** the mean best found value at each iteration, **B)** the mean Euclidean distance between the predicted optimum and the true optimum of the Hartmann surface. Lines denote mean values, shaded areas denote standard deviations.

6.3.4 Sensation Analysis

In order to investigate the effects of varying tACS current and frequency on participant reports of sensations during stimulation, linear mixed effect models were run with current and frequency as fixed factors and participant and block as nested random factors. This analysis showed an increase in participants' reports of itching, pain, burning and flashing lights in response to higher currents, as well as a decrease in burning, warmth and flashing lights in response to higher stimulation frequency. Full results of the analysis, including p-values corrected for multiple comparisons can be found in Table 6.1.

Variable	Predictors	Estimates	CI (95%)	df	t-value	p-value
Itching	Current	1.162	0.36–1.96	98	2.874	**0.012**
	Frequency	-0.015	-0.04– -0.02	98	-1.354	0.215
Pain	Current	0.553	0.10–0.55	98	2.400	**0.031**
	Frequency	-0.012	-0.03– -0.01	98	-1.963	0.079
Burning	Current	0.734	0.31–2.10	98	3.430	**0.004**
	Frequency	-0.018	-0.03 – -0.01	98	-3.053	**0.009**
Warmth	Current	0.334	-0.13–0.80	98	1.425	0.210
	Frequency	-0.023	-0.04– -0.01	98	-3.519	**0.004**
Fatigue	Current	-0.083	-0.86–0.69	98	-0.213	0.832
	Frequency	-0.008	-0.29–0.01	98	-0.742	0.501
Lights	Current	0.466	0.19–0.74	98	3.377	**0.004**
	Frequency	-0.010	-0.02– -0.00	98	-2.616	**0.021**

Table 6.1: **Results of the mixed effect model examining the effect of current and frequency on sensation data.** p-values are corrected using Benjamini-Hochberg corrections.

6.3.5 Baseline EEG analysis

The findings from the pBO model demonstrated differential effects of tACS according to baseline performance, with higher ability individuals predicted to re-

spond better to beta tACS (14-30Hz), whilst gamma tACS (>30) was shown to be beneficial for lower baseline individuals. Therefore, exploratory analysis was performed, examining the relationship between baseline ability and baseline beta power in frontal (F3) and parietal (P3) electrodes in the rs-EEG data. However, baseline ability was not found to be correlated with either frontal (r=-0.134, p=0.360) or parietal beta power (r=-0.118, p=0.421). Baseline Gamma power was not investigated as the higher frequencies are not reliably recorded by the equipment.

6.4 Discussion

In this chapter I have investigated the use of the pBO algorithm in selecting stimulation parameters for improving performance in an arithmetic task. Results of the study suggest that the pBO algorithm can tailor stimulation parameters to an individual's baseline ability in the arithmetic task. The BO algorithm iteratively identified more optimal tACS parameters, producing higher performance scores, suggesting successful performance of the algorithm.

Interestingly, the optimal stimulation parameters identified by the algorithm differed according to individuals' arithmetic baseline scores. This result suggests that the choice of the baseline score as the personalised variable was valid and that the algorithm has successfully characterised a portion of the inter-individual variation in the response to stimulation, potentially enabling the stimulation to achieve more reliable effects. However, whether this claim is accurate cannot be reliably determined in the present study and calls for a validation of the present results.

The differentiation in optimal stimulation parameters according to baseline

arithmetic performance does have interesting implications for the use of tACS in improving arithmetic abilities. It suggests that as one becomes more proficient at arithmetic, alternative cognitive processes, relying on different neural oscillatory frequency bands, are employed. An alternative interpretation of these results could be that tACS only benefits those individuals who are not as good at arithmetic, demonstrating a ceiling effect. This might mean that individuals who are more proficient at arithmetic do not benefit from tACS and so a stronger current or frequency might be more inclined to produce a negative effect, either through disrupting already efficient cognitive processes or simply acting as a distraction due to a higher chance of skin sensations. This would result in the BO algorithm identifying lower currents and frequencies as more optimal in higher ability individuals, as is the case in this experiment. It should be noted however, that the amount an individual could improve in response to stimulation did not appear to differ according to the predictions of the algorithm, regardless of their baseline ability, suggesting that tACS affected low and high baseline individuals equally. Although since the number of participants was heavily skewed to lower baseline abilities, no reliable statistical comparison can be made at this point, further requiring a validation study to be performed.

Simulation analysis demonstrated the clear benefit of the pBO algorithm over random sampling and the standard BO algorithm when there was a low level of noise in the data. The pBO algorithm not only converged on an optimum faster than random sampling and the BO algorithm, but it also identified the optima more accurately. As noise increased, the superiority of the pBO algorithm decreased, however it was never surpassed by the random sampling and BO algorithm. These results clearly demonstrate the superiority of the pBO algorithm in situations with

low levels of noise. whilst also illustrating that at higher levels of noise there is minimal benefit in using the algorithm over random sampling. However, the key component of the pBO algorithm is its ability to characterise the inter-individual variability aspect of this noise in the form of personalised variables. In this experiment, only one personalised variable was used in a modest sample size. At a larger scale, multiple personalised variables could be employed with larger numbers of participants, more evenly distributed across the range of those variables. If the personalised variables are carefully chosen to ensure they characterise the maximal amount of variability, as efficiently as possible, it may be possible to reduce the inter-individual noise, enabling the pBO algorithm to more accurately identify optimal stimulation parameters and further tailor the stimulation more closely to the individual.

Analysis of the sensation data collected in this experiment revealed similar findings to those in chapter 5, with higher current intensities increasing participants' ratings of a number of sensations including itching, pain and burning, whilst higher frequencies decreased sensations such as burning and warmth. The similarity between these results and those in chapter 5 is in spite of the use of different tES modalities. Although, it should be noted that both tACS and tRNS utilise alternating currents, and therefore may be altering neuronal function in a similar manner.

Taken together, the empirical and simulation results in this chapter strongly suggest that the pBO algorithm is capable of individualising stimulation parameters. However, as the results of chapters 3 and 4 have shown, there is a real need to validate the findings of BO algorithms when applied to complex human behaviour. Therefore, the following chapter will discuss a validation study of these results.

Chapter 7

Validating personalised Bayesian Optimisation for stimulation parameter selection

7.1 Introduction

In this chapter, I aim to validate the results found in chapter 6 using a within-subject experimental design where I compare the pBO-selected tACS parameters to a set of 'one-size-fits-all' standard parameters, as well as sham tACS. I hypothesised that not only would there be an improvement of arithmetic performance due to tACS, but that the pBO condition would improve performance compared to the standard parameters.

7.2 Methods

Twenty-six naive participants were recruited for the experiment. One participant failed to complete all experimental blocks, and was therefore discarded, leaving twenty-five participants (18 female, age (mean±sd) =19.52±2.18). The behavioural

task and arithmetic problems used in this chapter were identical to chapter 6. All participants gave informed consent and were compensated for their time. During the experiment, participants' resting-state EEG (rs-EEG) were recorded, and participants then performed ten practice trials of the arithmetic task. Participants then completed four blocks of fifty multiplication problems, the first of which was used to asses baseline arithmetic abilities as well as to normalise the remaining blocks, and the other three blocks were completed whilst receiving tACS. During these three tACS blocks, participants received either sham tACS, standard tACS or pBO-derived tACS, in a randomised order. rs-EEG was recorded after each stimulation block.

7.2.1 Stimulation

The alternating current stimulation was administered over the left frontoparietal network. The tACS was delivered via two stimulation (3.14 mm diameter) NG Pistim Ag/AgCl electrodes (F3 and P3) with one return electrode (Cz) using the Starstim 32 (Neuroelectrics, Barcelona). The impedances of the electrodes were held at <10 kΩ. During the tACS blocks, participants received either standard, pBO-derived or sham tACS. Standard tACS consisted of tACS delivered at 1mA and 6Hz, whilst pBO derived tACS parameters were automatically determined using the trained algorithm from chapter 6. Both standard and pBO tACS had a 30 second ramp up/down before and after the ten-minute stimulation block. Sham tACS consisted of only the ramp up, followed by immediate ramp down of stimulation. After completing a tACS block, subjects filled out a questionnaire in which they were asked several questions designed to gauge the level of sensation

experienced during stimulation (see Appendix A.2 for questionnaire).

7.2.2 EEG preprocessing and analysis

Resting state-EEG recordings were made at the start of the study (before baseline measurements) and immediately after every stimulation block. Electrophysiological data were obtained with eight gel Ag/AgCl electrodes (F3, P3, F4, P4, Fz, Cz, Pz, AF8) according to the international 10/20 EEG system using the Starstim 32 sensor system (Neuroelectrics, Barcelona, Spain), with no online filters. The ground consisted of adhesive active common mode sense (CMS) and passive driven right leg (DRL) electrodes which were positioned on the right mastoid. All EEG measurements had a duration of four minutes in which the subjects had their eyes open while watching a fixation point in the middle of the screen. Raw EEG data were recorded and used for offline analysis using EEGLAB (Delorme and Makeig 2004), in Matlab R2018b (Mathworks). Data were high-pass filtered at 0.1 Hz, and low-pass filtered at 50 Hz. Visual inspection was carried out to remove artefacts caused by muscle movement.

7.3 Results

7.3.1 Group Results

Outliers were removed if they were greater than 2 standard deviations from the mean, leading to the removal of 4 data points (2 standard tACS, 2 pBO). A repeated measures ANOVA on the group behavioural results demonstrated no effect of stimulation condition ($F(2,42)=1.373$, $p=.264$, $\eta_p^2=.061$), suggesting that

neither standard tACS or pBO tACS significantly improved arithmetic performance compared to sham tACS (Figure 7.1).

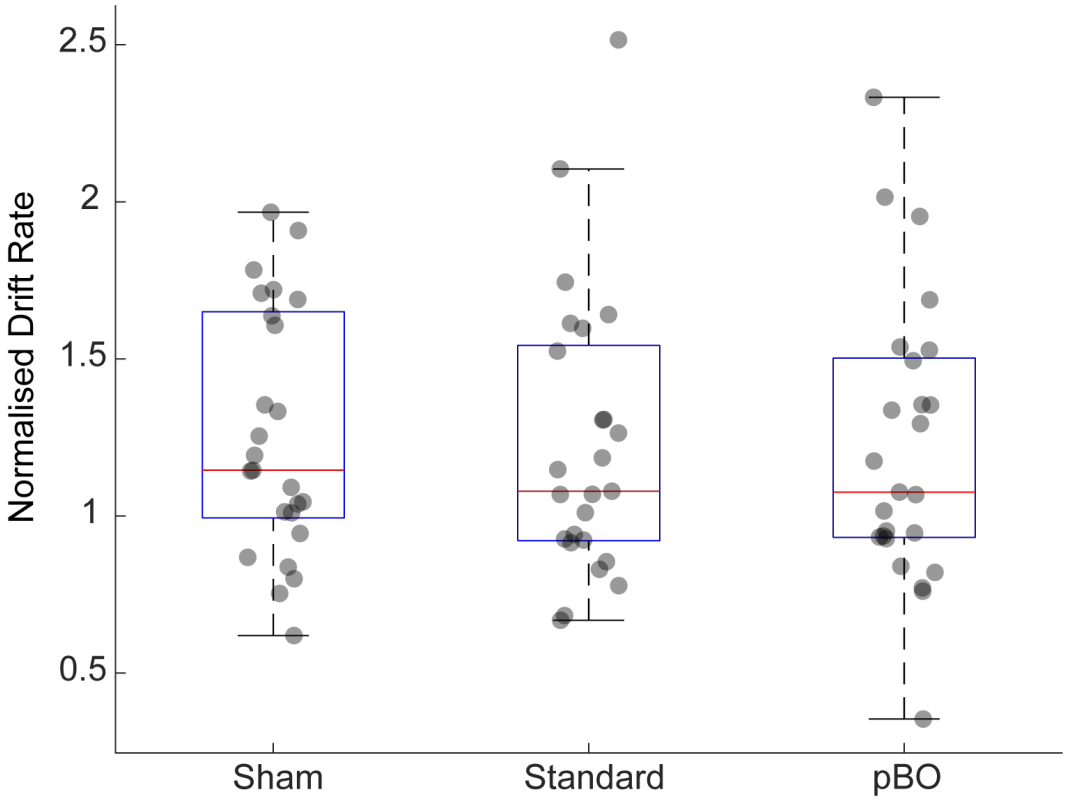


Figure 7.1: Group results show no difference between any of the three groups.

Results from chapter 6 indicated a possible difference in optimal stimulation parameters according to baseline ability. As this difference in optimal parameters could potentially enable individuals with different baseline abilities to perform at the same level, a linear regression was performed investigating the effects of stimulation type and baseline ability on performance in the pBO and standard tACS conditions. Of particular interest were any possible interactions between stimulation type and baseline ability, potentially with lower baseline individuals performing better with receiving pBO-derived stimulation parameters than when receiving

'standard' parameters. However, the results of the regression showed no significant interaction (β=0.513 t(3,42)=0.847, p=0.402) or any main effects of stimulation type (β=0.425, t(3,42)=1.059, p=0.296), or baseline (β=0.472, t(3,42)=0.984, p=0.331).

7.3.2 pBO analysis

A visual comparison of the participants' arithmetic performance when receiving pBO-derived tACS and the predictions from the pBO model demonstrated that the pBO algorithm overestimated the effects of stimulation on arithmetic performance (Figure 7.2).

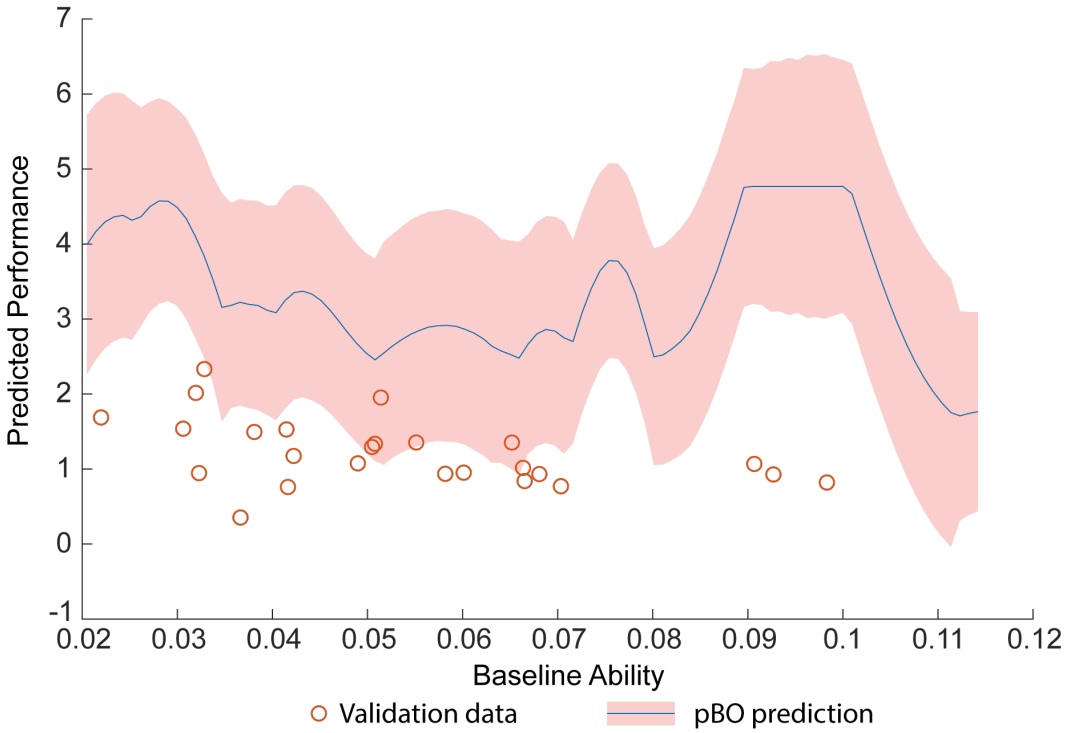


Figure 7.2: A comparison of the performance predictions from the pBO model vs the new data gathered in the validation experiment.

7.3.3 Control Analysis

Linear regression was performed to test for any effects of the order in which participants received the three different stimulation conditions, as well as any effect of baseline ability. No effect of order (β=0.251, t(2,72)=1.7, p=0.096) or baseline ability (β=0.066, t(2,72)=0.449, p=0.655) were found.

Further control analysis aimed to determine whether participants were appropriately blinded. A Chi-squared test was performed to test whether the occurrence of participants reporting that they thought they received active tACS differed in the three groups. A significant difference was found (X^2(2,n=71)=10.98, p=0.004, V=0.393), with 13 participants guessing that they received stimulation in the sham condition, 6 in the standard tACS condition, and 2 in the pBO tACS condition. These results indicate that blinding was not effective in disguising the different stimulation conditions; however the effect was due to a decrease in stimulation guesses in the active stimulation conditions. To investigate whether this difference in perceived stimulation condition affected an individual's performance, a mixed effect model was performed examining the effects of stimulation condition and perceived condition on performance, with random effects for subject. A negative effect of standard tACS was found on participant performance compared to sham (β=-0.206, t(41)=-2.836, p=0.007), but no main effect of BO-derived tACS (β=-0.072, t(41)=-1.017, p=0.315). There was no main effect of perceived stimulation condition (β=-0.145, t(41)=-1.644, p=0.108), however perceived condition did interact with both standard tACS (β=0.322, t(41)=2.414, p=0.020) and BO-derived tACS (β=-0.477, t(41)=-2.484, p=0.017). As can be seen in Figure 7.3, perception of active stimulation slightly improved performance in the standard tACS group,

whereas it considerably reduced it in the BO-derived tACS group. It should be noted however, that only 2 participants guessed that they received real stimulation in the BO-derived tACS group.

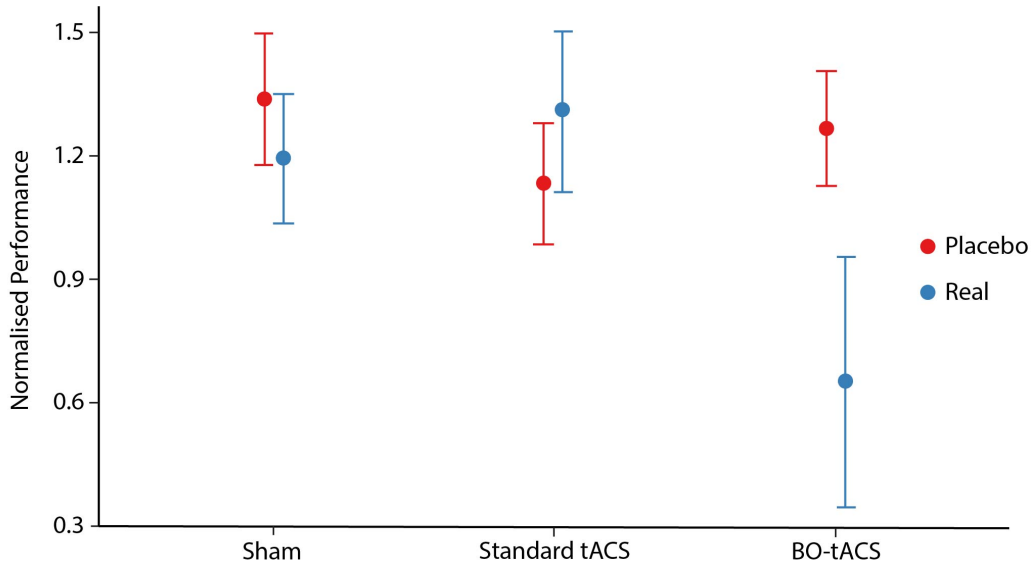


Figure 7.3: **Interactions between perceived and actual stimulation condition affecting performance.** Colour indicates whether participants guessed that they received real or placebo tACS.

To further investigate the findings from the blinding analysis, the effect of condition was examined on the sensation data gathered in this experiment. No differences were found for ratings of itching (X^2(10,n=71)=13.64, p=0.190, V=0.310), pain (X^2(10,n=71)=6.617, p=0.761, V=0.216), burning (X^2(10,n=71)=11.26, p=0.338, V=0.282), warmth (X^2(6,n=71)=5.702, p=0.457, V=0.200), fatigue (X^2(8,n=71)=8.792, p=0.360, V=0.249) or reports of flashing lights(X^2(4,n=71)=3.511, p=0.476, V=0.157).

7.3.4 EEG Analysis

Results from the pBO model in chapter 6 suggested a differential effect of tACS according to baseline performance, with higher ability individuals predicted to respond better to beta tACS (14-30Hz). EEG analysis in chapter 6 found no correlation between baseline ability and beta power. These results were validated in the current chapter, as baseline ability was not found to be correlated with either frontal (r=0.122, p=0.578) or parietal beta power (r=0.111, p=0.614) in the rs-EEG recorded prior to the baseline behavioural block.

7.4 Discussion

In this chapter I aimed to validate the results of chapter 6 by comparing the pBO derived tACS parameters to a 'standard' parameter combination and sham tACS. I hypothesised that there would be an improvement in arithmetic performance due to tACS and that the pBO condition would improve performance compared to the standard parameters. However, the results clearly show no effect of either pBO-derived tACS or 'standard' tACS when compared to sham, suggesting that stimulation did not improve arithmetic performance. Additionally, no interaction was found between baseline ability and stimulation type when investigating the pBO and standard stimulation conditions, suggesting that the pBO-derived tACS parameters do not improve performance in lower baseline individuals to the level of higher baseline individuals, a result that makes sense if there is indeed no effect of stimulation on this task. Whilst these results are unfortunate in that no effect of tACS was found, they are not surprising as minimal work has been done utilising tACS to improve arithmetic performance. Previous work has implicated the left

frontalparietal network in arithmetic processing (Ischebeck et al. 2007; Klein et al. 2016; Matejko and Ansari 2015; Nieder 2016) and has shown that most arithmetic-related frequency band changes occur in the theta and alpha bands (Grabner and De Smedt 2012; Hinault and Lemaire 2016). These results suggested that tACS delivered over the left frontalparietal network might be able to modulate arithmetic performance, however this was not the case. This may be because the pBO algorithm was not successful in identifying the optimal stimulation parameters for improving arithmetic performance or that there was too much noise in the current validation study in order to detect an effect. Alternatively it may be that tACS does not improve arithmetic performance at all, or just at this electrode location. Given the previous literature highlighting the importance of the frontal brain regions (Artemenko et al. 2015), it may be that stimulating more frontally would be more beneficial. Electrode location is a parameter that could also be included in the pBO algorithm in future research, however in the current research it was not included for simplicity and to keep the required number of participants down.

A further visual comparison of the results compared the data collected in this experiment in the pBO stimulation condition, and the predictions of participant performance made by the pBO model in chapter 6. This comparison demonstrated that the pBO model had overestimated participant performance in response to stimulation. This is likely due to the large levels of variability in participant responses causing the model to believe a certain set of stimulation parameters are particularly beneficial. Unfortunately, the addition of the personalised variable means the model can only re-sample around peaks in the search space when it is presented with a participant with a suitable baseline ability, therefore peaks caused by random variability in participant performance cannot be retested and are therefore

inaccurately modelled. This explains why predictions of normalised performance are slightly higher for participants with higher baselines, as the sample in chapter 6 was skewed to the lower end of the baseline range, therefore not allowing the pBO algorithm to retest peaks in the experimental space. Unfortunately, the only solution for this problem is to employ larger sample sizes in pBO studies and to ensure an even distribution of baseline abilities across the spectrum. The exact sample size required for a pBO algorithm is unknown. In chapter 6, 50 participants were tested, producing approximately 50 data points per dimension explored by the pBO algorithm (50D), double the 25D that a recent study used in a noisy environment (Letham et al. 2019). Clearly the exact sample size required depends on the expected effect size, the general shape of the search space (multiple optima will increase the sample size required) and how many dimensions are being searched. Future work utilising pBO in noisy environments should aim for large sample sizes, possibly through a multi-centre study, therefore enabling the collection of a wider distribution of participants' personalised variables.

Further analysis investigated the effectiveness of the blinding in the paradigm. Results showed that when participants thought they were receiving active stimulation, there was no behavioural difference in the sham group. However, in the standard tACS group, it slightly increased performance, whilst in the BO-tACS group it very clearly decreased performance. Whilst these are interesting results which suggest an interacting effect between stimulation group and perception of stimulation on behavioural effects, it should be noted that only two participants in the BO-tACS group thought they were receiving active stimulation, calling the validity of these results into question and limiting the conclusions that can be drawn.

Chapter 8

Discussion

8.1 Summary

The research presented in this thesis investigated how the use of BO could improve the outcomes of neurointerventions. Specifically, I focused on improving the effectiveness of tES in modulating behaviour by utilising the BO algorithm to select stimulation parameters. The key aims of this thesis were:

- To apply BO to parameter selection of tES in order to identify an optimal set of stimulation parameters for improving performance on a cognitive task. This was done in chapter 3.
- To attempt to reduce the variation in response to tES by personalising the stimulation parameters to the individual. This was done in chapters 5 and 6.
- To thoroughly validate all results of the BO algorithm to ensure reliability of the selected tES parameters. This was done in chapters 4 and 7.

In **chapter 3** I demonstrated the first attempt at using BO with tES in a complex cognitive task. In this study I built on work by Lorenz, Simmons, et

al. (2019) and Cole, Bagic, et al. (2010) to utilise BO to search for the optimal stimulation current intensity and frequency of tRNS to improve transfer in RITL in a between-subject design. The results of this chapter were promising as they suggested that the algorithm was able to identify optimal parameters even when tES is known to have large levels of variability in its effects. Further, the selected optimal tRNS parameters supported mechanistic theories of tRNS, namely ripple oscillations, as the selected frequency was within the range expected in that theory (Moliadze, Antal, and Paulus 2010). However, simulation analysis of these results, whilst demonstrating the superiority of BO over a random search algorithm, showed that a very large effect size was needed to reliably identify an optimum using the algorithm. As both complex behaviours such as RITL and tES techniques can have a large amount of variation, leading to increased noise for the BO algorithm, these results clearly demonstrated the need to validate these findings, a step which previous research using BO in human behaviour has not always done.

In **chapter 4** I ran a well-powered study of 80 participants, comparing the stimulation parameters determined by the BO algorithm to a sham stimulation condition. The BO-derived parameters were not found to be more effective at improving transfer in RITL than the sham tRNS. This experiment demonstrated the importance of validating the results of BO algorithms in tES.

In **chapter 5** I built upon work by A. Krause and Ong (2011) and my work in chapters 3 and 4 and demonstrated the first use of a pBO algorithm capable of incorporating personalised variables in order to tailor the stimulation parameters to the individual. In this chapter I aimed to determine the optimal tACS current intensity and frequency for improving sustained attention whilst incorporating participant TBR. Unfortunately, the pBO algorithm failed to find any trends in

the behavioural data in response to different stimulation parameters and resorted to random sampling of the parameter combinations. These results indicated that TBR was either too complicated or noisy a personalised variable to be easily used in the pBO algorithm, or did not account for a sufficient amount of intra-subject variability in performance to allow stimulation parameters to vary across its range.

In **chapter 6** I demonstrated the use of a pBO algorithm to incorporate personalised variables in the selection of tRNS parameters to improve arithmetic abilities. Importantly, a more simplistic personalised variable was chosen in comparison to chapter 5, in the form of baseline arithmetic ability. Results from this experiment were promising in that the pBO algorithm not only identified parameter combinations that led to greater improvement in performance, but also identified different tRNS parameters depending on participants' baseline ability. Importantly, simulation analysis demonstrated that the pBO algorithm was more effective at identifying an optimum than both standard BO and random search algorithms, although this difference was limited at higher levels of noise.

In **chapter 7** I attempted to validate the results of chapter 6 by comparing the pBO-derived tRNS parameters with a 'standard' set of parameters as well as sham tRNS. There was no difference between pBO-derived, 'standard' and sham tRNS parameters. Not only did this again highlight the need for validation in BO-tES experiments, but it also called into question the effectiveness of BO and pBO algorithms in tES experiments in sample sizes <100 participants.

In the following sections, I will discuss the advances made, and the lessons learnt, from this thesis, surrounding BO and its potential application to neurointerventions. I will also discuss the limitations of BO and my work, and outline important directions for future work.

8.2 Advances and Limitations

8.2.1 Data-driven tES parameter selection

The application of BO to provide a data-driven method of selecting tES parameters is a relatively untested one when applied to complex cognitive behaviours. Previous research has used BO either in complex cognitive behaviours (Lorenz, Violante, et al. 2018) or to select tES parameters in physiological responses (Lorenz, Simmons, et al. 2019). Specifically, the ability of the algorithm to balance exploration of new parameter combinations with the exploitation of those it already believes to be beneficial is a clear improvement over the traditional method of comparing a small number of different parameters. However, this method has its downsides in that whilst it covers a large number of different parameter combinations, its power is in detecting trends across the entire experimental space. If the optimal parameter combination is situated on a relatively narrow peak in the experimental space, where there are multiple optima present in the space, then the algorithm may not converge on it. Alternatively, in situations in which there is a large amount of noise, the algorithm does not have enough power to detect the underlying trends (Rasmussen and Williams 2006), as evidenced by simulation analysis in chapter 3. Therefore, researchers must exercise prudence in determining when to deploy this algorithm.

The use of a between-subject design to select a one-size-fits-all set of parameters is a novel feature of this thesis. As shown by chapter 4, it is likely that in the case of tES, a one-size-fits-all set of parameters is not optimal given the levels of variation previously observed in response to stimulation. However, this technique does have potential applicability to other areas of neurointervention in which there is less

intra-individual variation.

8.2.2 Personalised tES parameters

The decision in chapters 5 and 6 to expand the BO algorithm to include personalised variables is an advancement unique to this work in the domain of human cognition, having previously been used in the machine learning literature (A. Krause and Ong 2011). This work highlights the need for large sample sizes as well as the careful choice of the personalised variable. Specifically, the personalised variable must account for as large an amount of the variability between individuals as is possible, potentially through the use of multiple personalised variables in larger studies. If future work is able to successfully utilise this method to tailor neurointervention, it promises to be a simple way to personalise the intervention. This would be especially promising for clinical use where repeated testing of an individual to determine treatment is not advised. However, the clear limitation of this method is that a large initial study must be conducted to allow the pBO algorithm to identify the optimal parameters.

Another limitation of the pBO approach used in this work is that in order to collect a large amount of data, participants were exposed to multiple stimulation parameter combinations over the course of one session. Whilst this enabled multiple data points from each participant to be collected, a feature which is conducive to convergence of the pBO algorithm, this ignores any carry-over effects of stimulation from one parameter combination to the next. Studies investigating the lasting effects of stimulation have found that at very short durations of tACS, there are no lasting neuronal effects (**struber**), however delivery of tACS over a longer du-

rations (Clancy et al. 2018; Moliadze, Sierau, et al. 2019; Heise et al. 2019) or over repeated sessions has been shown to have more lasting effects(Fresnoza et al. 2018). Similarly, both single session (Terney et al. 2008) and repeated session (Snowball et al. 2013) paradigms of tRNS have been shown to have continued effects after the cessation of stimulation. The potential presence of carry over effects is a major confound in the work in this thesis, as it would make it more difficult to disentangle the true effect of a particular set of stimulation parameters as the lasting effects of those delivered previously may still be being detected. This would cause significant problems for the pBO algorithms in chapters 5 and 6 as carry over effects would make the data more noisy and therefore make it harder to identify an optimum set of parameters. Ideally, separate sessions would be used to collect data for each parameter combination an individual was exposed to; however this is impractical due to the amount of time involved, therefore I opted to administer different stimulation parameters in one session. Whilst this does add noise to the measures, it was hoped that the pBO algorithm would be able to deal with the added noise and still identify optimal parameter combinations as the algorithm is able to retest any point in the search space. Future work should seek to counteract the presence of carry over effects by either conducting a separate session for each parameter combination, or by somehow weighting the recorded data in terms of the order in which it was collected. For example, a data point collected as the first parameter combination used in that session would influence the GP's surrogate model more than one that was collected second in the experimental session. Therefore it would be possible to minimise the effect of carry over effects whilst still enabling the collection of multiple data points in one session.

8.2.3 Validation of BO results

An important feature of this thesis is that all of the promising results from the BO algorithm were validated using a separate study with good experimental power. As the stimulation parameters identified by the BO algorithms were not found to be better than sham stimulation, this step was shown to be very important. Alternative methods of validating the results of BO algorithms include replicating the results of the BO algorithm in other participants, if using a within-subject design, or through simulation analysis (Lorenz, Simmons, et al. 2019). Whilst I have conducted simulation analysis demonstrating the abilities of BO algorithms on simulated data, I also showed that these are not adequate substitutes for an experimental validation, as no simulation fully captures true human responses or the effects of stimulation.

8.3 Future work

8.3.1 Study design

As seen in the work in this thesis, results of the BO algorithms, whilst initially appearing positive, can be unreliable and fail to be validated. Although one of the benefits of BO is that a smaller sample size is needed than if each parameter combination were to be tested many times, in order to cope with the variability inherent in human cognition, larger samples than the ones used in this thesis are advisable. Ideally, large multi-centre studies should be run, enabling the collection of large amounts of data. The addition of more data also means that multiple personalised variables would be able to be used, potentially accounting for more

of the variability between individuals in response to stimulation. The benefit of this work would be that once it is complete, tES parameters could be tailored to any new individuals requiring stimulation with a high degree of personalisation and with minimal effort.

8.3.2 Adapting to variability in behavioural responses

A potential source of inefficiency in the BO algorithms employed in this thesis and in those used in human research previously, is that when the GP models the variability within the objective function, it does so in a uniform manner across the experimental space. This can be inefficient and is an inaccurate representation of the objective function for some parameters. For example, when utilising a personalised variable in the pBO algorithm consisting of baseline ability, it may be expected that individuals with naturally higher ability will be more consistent in their behavioural performance than those at lower end of the range. Properly modelling the variability, and allowing it to differ across the range of the experimental space, can enable the algorithm to be more efficient as less samples are needed in areas where there is lower variability and more samples needed in areas of high variability.

8.3.3 Hierarchical BO

The BO algorithms presented in this thesis utilised data collected from multiple participants. In chapter 3 only one data point was collected per participant, however in chapters 5 and 6, multiple data points were collected per participant. In the models used in this thesis there was no grouping of data points from the

same participant, potentially increasing noise in the model. In order to negate this, hierarchical BO could be implemented to group data from each participant, therefore allowing for the objective function to be modelled more efficiently (Hensman, Lawrence, and Rattray 2013; Ouyang et al. 2020). An important development in this field would be to create acquisition functions capable of capturing the different levels of inter- and intra-individual variability whilst allowing the algorithm to optimise across participants. This is particularly important when applying these techniques to situations in which there is a small effect size such as tES.

8.3.4 Statistical improvements

An important question when utilising BO is how to perform statistical inference on the results to determine the likelihood of identifying a real underlying objective function instead of a random one. In this thesis, the statistics performed on the trained BO algorithm and GP surrogate model have been extremely limited as the range of statistical techniques applicable to the situation is very limited. In order to appropriately gauge the performance of BO algorithms in human research, novel statistical techniques must be developed, which will be particularly relevant in areas in which there is a high degree of variability.

8.4 Concluding remarks

This thesis has investigated the use of BO in parameter selection for tES, and has expanded upon previous work by attempting to personalise the results to further improve behavioural outcomes and reduce variability in response to tES. Whilst my results have not been consistently positive, I believe that further work in this

area can enable this approach to become common practice in research and clinical settings alike. In particular, if pBO can be appropriately utilised then it promises to further inform future tES research in maximising behavioural and clinical effects whilst minimising costs such as safety risks and participant time. In summary, I argue that BO and pBO are important new tools in parameter selection for neurointerventions.

Bibliography

Ammann, C., M. A. Lindquist, and P. A. Celnik (2017). "Response variability of different anodal transcranial direct current stimulation intensities across multiple sessions". In: *Brain Stimulation* 10.4, pp. 757–763.

Antal, A., L. Chaieb, et al. (2014). "BDNF gene polymorphisms and motor cortical plasticity in healthy humans: when should we consider it". In: 9091.2, pp. 1–13.

Antal, A., D. Keeser, et al. (July 2015). "Conceptual and Procedural Shortcomings of the Systematic Review "Evidence That Transcranial Direct Current Stimulation (tDCS) Generates Little-to-no Reliable Neurophysiologic Effect Beyond MEP Amplitude Modulation in Healthy Human Subjects: A Systematic R". In: *Brain Stimulation* 8.4, pp. 846–849.

Arden, R. et al. (Dec. 2010). "Neuroimaging creativity: A psychometric view". In: *Behavioural Brain Research* 214.2, pp. 143–156.

Arns, M., C. K. Conners, and H. C. Kraemer (2013). "A Decade of EEG Theta/Beta Ratio Research in ADHD: A Meta-Analysis". In: *Journal of Attention Disorders* 17.5, pp. 374–383.

Artemenko, C. et al. (2015). "Differential influences of unilateral tDCS over the intraparietal cortex on numerical cognition". In: *Frontiers in Human Neuroscience* 9.March, pp. 1–8.

Bartrés-Faz, D., S. Rossi, and A. Pascual-Leone (2016). "Editorial: Non-invasive brain stimulation and plasticity changes in aging". In: *Frontiers in Aging Neuroscience* 8.MAY, pp. 8–10.

Batsikadze, G. et al. (2013). "Partially non-linear stimulation intensity-dependent effects of direct current stimulation on motor cortex excitability in humans". In: *Journal of Physiology* 591.7, pp. 1987–2000.

Battleday, R. M. et al. (2014). "Mapping the mechanisms of transcranial alternating current stimulation: a pathway from network effects to cognition." In: *Frontiers in psychiatry* 5, p. 162.

Benwell, C. S. et al. (2019). "Frequency and power of human alpha oscillations drift systematically with time-on-task". In: *NeuroImage* 192.March, pp. 101–114.

Bereznicki, H. G. et al. (2019). "Assessment of double blinding in tES research: A call for the establishment of standard procedures". In: *Brain Stimulation* 12.6, pp. 1608–1609.

Bikson, M., A. R. Brunoni, et al. (2018). "Rigor and reproducibility in research with transcranial electrical stimulation: An NIMH-sponsored workshop." In: *Brain stimulation* 11.3, pp. 465–480.

Bikson, M., Z. Esmaeilpour, et al. (2019). "Transcranial electrical stimulation nomenclature". In: *Brain Stimulation* 12.6, pp. 1349–1366.

Brem, A. K. et al. (Sept. 2018). "Modulating fluid intelligence performance through combined cognitive training and brain stimulation". In: *Neuropsychologia* 118, pp. 107–114.

Brochu, E., V. M. Cora, and N. D. Freitas (2010). "A Tutorial on Bayesian Optimization of Expensive Cost Functions , with Application to Active User Modeling and Hierarchical Reinforcement Learning". In: arXiv: `arXiv:1012.2599v1`.

Button, K. S. et al. (2013). "Power failure: Why small sample size undermines the reliability of neuroscience". In: *Nature Reviews Neuroscience* 14.5, pp. 365–376.

Buzsáki, G. and F. L. da Silva (2012). "High frequency oscillations in the intact brain". In: *Progress in Neurobiology* 98.3, pp. 241–249.

Chaieb, L., A. Antal, and W. Paulus (2008). "Gender-specific modulation of short-term neuroplasticity in the visual cortex induced by transcranial direct current stimulation". In: *Visual Neuroscience* 25.1, pp. 77–81.

Chew, T., K. A. Ho, and C. K. Loo (2015). "Inter- and intra-individual variability in response to transcranial direct current stimulation (tDCS) at varying current intensities". In: *Brain Stimulation* 8.6, pp. 1130–1137.

Clancy, K. J. et al. (2018). "Lasting connectivity increase and anxiety reduction via transcranial alternating current stimulation". In: *Social Cognitive and Affective Neuroscience* 13.12, pp. 1305–1316.

Clayton, M. S., N. Yeung, and R. Cohen Kadosh (2019). "Electrical stimulation of alpha oscillations stabilises performance on visual attention tasks". In: *Journal of Experimental Psychology: General* 142.2, pp. 203–220.

Cohen Kadosh, R., S. Soskic, et al. (2010). "Modulating neuronal activity produces specific and long-lasting changes in numerical competence". In: *Current Biology* 20.22, pp. 2016–2020.

Cohen Kadosh, R. and V. Walsh (2007). "Dyscalculia". In: *Current Biology* 17.22, pp. 946–947.

Cole, M. W., A. Bagic, et al. (2010). "Prefrontal Dynamics Underlying Rapid Instructed Task Learning Reverse with Practice". In: *Journal of Neuroscience* 30.42, pp. 14245–14254.

Cole, M. W., T. S. Braver, and N. Meiran (2017). "The task novelty paradox: Flexible control of inflexible neural pathways during rapid instructed task learning". In: *Neuroscience and Biobehavioral Reviews* 81, pp. 4–15.

Cole, M. W., J. A. Etzel, et al. (2011). "Rapid Transfer of Abstract Rules to Novel Contexts in Human Lateral Prefrontal Cortex". In: *Frontiers in Human Neuroscience* 5.November, pp. 1–13. arXiv: NIHMS150003.

Cole, M. W., L. M. Patrick, et al. (2018). "A role for proactive control in rapid instructed task learning". In: *Acta Psychologica* 184.November 2016, pp. 20–30.

Cole, M. W., J. R. Reynolds, et al. (2013). "Multi-task connectivity reveals flexible hubs for adaptive task control". In: *Nature Neuroscience* 16.9, pp. 1348–1355.

Costa, T. L. et al. (May 2015). "Contrasting effects of transcranial direct current stimulation on central and peripheral visual fields". In: *Experimental Brain Research* 233.5, pp. 1391–1397.

Davis, N. J. (2020). "Variance in cortical depth across the brain surface: Implications for transcranial stimulation of the brain". In: *European Journal of Neuroscience* August, pp. 1–12.

Delorme, A. and S. Makeig (2004). "EEGLAB: an open source tollbox for analysis of single-trial EEG dynamics including independent component analysis". In: *Journal of Neuroscience Methods* 134.1, pp. 9–21.

Durand, A. et al. (2018). "Contextual Bandits for Adapting Treatment in a Mouse Model of de Novo Carcinogenesis". In: *Mlhc-2018* 85, pp. 1–15.

Ego-Stengel, V. and M. A. Wilson (2010). "Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat". In: *Hippocampus* 20.1, pp. 1–10.

Esmaeilpour, Z. et al. (Mar. 2018). "Incomplete evidence that increasing current intensity of tDCS boosts outcomes". In: *Brain Stimulation* 11.2, pp. 310–321.

Fertonani, A., C. Pirulli, and C. Miniussi (2011). "Random Noise Stimulation Improves Neuroplasticity in Perceptual Learning". In: *The Journal of Neuroscience* 31.43, pp. 15416–15423.

Fertonani, A., S. Rosini, et al. (2010). "Naming facilitation induced by transcranial direct current stimulation". In: *Behavioural Brain Research* 208.2, pp. 311–318.

Feurra, M. et al. (2013). "State-dependent effects of transcranial oscillatory currents on the motor system: what you think matters." In: *The Journal of Neuroscience* 33.44, pp. 17483–17489.

Fonteneau, C. et al. (2019). "Sham tDCS: A hidden source of variability? Reflections for further blinded, controlled trials". In: *Brain Stimulation* 12.3, pp. 668–673.

Fortenbaugh, F. C. et al. (2018). "Tracking behavioural and neural fluctuations during sustained attention: A robust replication and extension". In: *NeuroImage* 171, pp. 148–164.

Fresnoza, S. et al. (2018). "The effects of transcranial alternating current stimulation (tACS) at individual alpha peak frequency (iAPF) on motor cortex excitability in young and elderly adults". In: *Experimental Brain Research* 236.10, pp. 2573–2588.

Fröhlich, F. and D. A. McCormick (2010). "Endogenous electric fields may guide neocortical network activity". In: *Neuron* 67.1, pp. 129–143.

Fumagalli, M. et al. (2010). "Brain switches utilitarian behavior: Does gender make the difference?" In: *PLoS ONE* 5.1.

Girardeau, G. et al. (2009). "Selective suppression of hippocampal ripples impairs spatial memory". In: *Nature Neuroscience* 12.10, pp. 1222–1223.

Grabner, R. H. and B. De Smedt (2012). "Oscillatory EEG correlates of arithmetic strategies: a training study." eng. In: *Frontiers in psychology* 3, p. 428.

Greinacher, R. et al. (Oct. 2019). "The time course of ineffective sham-blinding during low-intensity (1 mA) transcranial direct current stimulation". In: *European Journal of Neuroscience* 50.8, pp. 3380–3388.

Groen, O. van der and N. Wenderoth (2016). "Transcranial Random Noise Stimulation of Visual Cortex: Stochastic Resonance Enhances Central Mechanisms of Perception". In: *The Journal of Neuroscience* 36.19, pp. 5289–5298. arXiv: NIHMS150003.

Guerra, A., V. López-Alonso, et al. (2020). "Variability in non-invasive brain stimulation studies: Reasons and results". In: *Neuroscience Letters* 719.December 2017, p. 133330.

Guerra, A., A. Pogosyan, et al. (Oct. 2016). "Phase Dependency of the Human Primary Motor Cortex and Cholinergic Inhibition Cancelation During Beta tACS." eng. In: *Cerebral cortex (New York, N.Y. : 1991)* 26.10, pp. 3977–3990.

Hadida, J. et al. (2018). "Bayesian Optimisation of Large-Scale Biophysical Networks". In: *NeuroImage* 174.February, pp. 219–236.

Harty, S. and R. Cohen Kadosh (2019). "Suboptimal Engagement of High-Level Cortical Regions Predicts Random-Noise-Related Gains in Sustained Attention". In: *Psychological Science* 30.9, pp. 1318–1332.

Heise, K. F. et al. (2019). "Distinct online and offline effects of alpha and beta transcranial alternating current stimulation (tACS) on continuous bimanual performance and task-set switching". In: *Scientific Reports* 9.1, pp. 1–16.

Hensman, J., N. D. Lawrence, and M. Rattray (2013). "Hierarchical Bayesian modelling of gene expression time series across irregularly sampled replicates and clusters". In: *BMC Bioinformatics* 14.1.

Hernández-Lobato, J. M., M. W. Hoffman, and Z. Ghahramani (2014). "Predictive entropy search for efficient global optimization of black-box functions". In: *Advances in Neural Information Processing Systems* 1.January, pp. 918–926. arXiv: 1406.2541.

Herrmann, C. S. et al. (2013). "Transcranial alternating current stimulation : a review of the underlying mechanisms and modulation of cognitive processes". In: *Frontiers in human neuroscience* 7.279, pp. 1–13.

Hinault, T. and P. Lemaire (Aug. 2016). "What does EEG tell us about arithmetic strategies? A review." eng. In: *International journal of psychophysiology : official journal of the International Organization of Psychophysiology* 106, pp. 115–126.

Horvath, J. C. (May 2015). "New Quantitative Analyses Following Price & Hamilton's Critique do not Change Original Findings of Horvath et al." In: *Brain Stimulation* 8.3, pp. 665–666.

Horvath, J. C., O. Carter, and J. D. Forte (Oct. 2016). "No significant effect of transcranial direct current stimulation (tDCS) found on simple motor reaction time comparing 15 different simulation protocols". In: *Neuropsychologia* 91, pp. 544–552.

Hu, Z. et al. (2019). "Fraud regulating policy for E-Commerce via constrained contextual bandits". In: *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS* 3, pp. 1377–1385.

Inghilleri, M. et al. (May 2004). "Ovarian hormones and cortical excitability. An rTMS study in humans." eng. In: *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology* 115.5, pp. 1063–1068.

Ischebeck, A. et al. (2007). "Imaging early practice effects in arithmetic". In: *NeuroImage* 36, pp. 993–1003.

Jacobs, J. et al. (2006). "EEG oscillations and recognition memory : Theta correlates of memory retrieval and decision making". In: *NeuroImage* 32, pp. 978–987.

Johnson, L. et al. (2020). "Dose-dependent effects of transcranial alternating current stimulation on spike timing in awake nonhuman primates". In: *Science Advances* 6.36.

Joy, T. T. et al. (2016). "Hyperparameter tuning for big data using Bayesian optimisation". In: *Proceedings - International Conference on Pattern Recognition* 0, pp. 2574–2579.

Kanai, R. et al. (Dec. 2008). "Frequency-Dependent Electrical Stimulation of the Visual Cortex". In: *Current Biology* 18.23, pp. 1839–1843.

Kaufmann, L. et al. (2013). "Dyscalculia from a developmental and differential perspective". In: *Frontiers in Psychology* 4.AUG, pp. 1–5.

Klein, E. et al. (Mar. 2016). "Considering structural connectivity in the triple code model of numerical cognition: differential connectivity for magnitude processing and arithmetic facts." eng. In: *Brain structure & function* 221.2, pp. 979–995.

Koenigs, M. et al. (Jan. 2009). "Bilateral frontal transcranial direct current stimulation: Failure to replicate classic findings in healthy subjects." In: *Clinical*

neurophysiology : official journal of the International Federation of Clinical Neurophysiology 120.1, pp. 80–4.

Krause, A. and C. S. Ong (2011). "Contextual Gaussian Process Bandit Optimization". In: *Advances in Neural Information Processing Systems*, pp. 2447–2455.

Krause, B. and R. Cohen Kadosh (2014). "Not all brains are created equal: the relevance of individual differences in responsiveness to transcranial electrical stimulation." In: *Frontiers in systems neuroscience* 8.February, p. 25.

Lafon, B. et al. (Oct. 2017). "Low frequency transcranial electrical stimulation does not entrain sleep rhythms measured by human intracranial recordings." eng. In: *Nature communications* 8.1, p. 1199.

Learmonth, G. et al. (2017). "No Interaction between tDCS Current Strength and Baseline Performance: A Conceptual Replication." In: *Frontiers in neuroscience* 11, p. 664.

Letham, B. et al. (2019). "Constrained Bayesian optimization with noisy experiments". In: *Bayesian Analysis* 14.2, pp. 495–519. arXiv: 1706.07094.

Liu, A. et al. (2018). "Immediate neurophysiological effects of transcranial electrical stimulation". In: *Nature Communications* 9.1.

Looi, C. Y. et al. (2016). "Combining brain stimulation and video game to promote long-term transfer of learning and cognitive enhancement". In: *Scientific Reports* 6.August 2015, pp. 1–12.

López-Alonso, V. et al. (2015). "Intra-individual variability in the response to anodal transcranial direct current stimulation". In: *Clinical Neurophysiology* 126.12, pp. 2342–2347.

Lorenz, R., R. P. Monti, et al. (2016). "Towards tailoring non-invasive brain stimulation using real-time fMRI and Bayesian optimization". In: *2016 International Workshop on Pattern Recognition in Neuroimaging (PRNI)*, pp. 1–4. arXiv: 1605.01270.

Lorenz, R., L. E. Simmons, et al. (Nov. 2019). "Efficiently searching through large tACS parameter spaces using closed-loop Bayesian optimization". In: *Brain Stimulation* 12.6, pp. 1484–1489.

Lorenz, R., I. R. Violante, et al. (2018). "Dissociating frontoparietal brain networks with neuroadaptive Bayesian optimization". In: *Nature Communications* 9.1.

Matejko, A. A. and D. Ansari (Jan. 2015). "Drawing connections between white matter and numerical and mathematical cognition: a literature review." eng. In: *Neuroscience and biobehavioral reviews* 48, pp. 35–52.

Molero-Chamizo, A. et al. (Feb. 2018). "Poststimulation time interval-dependent effects of motor cortex anodal tDCS on reaction-time task performance". In: *Cognitive, Affective, & Behavioral Neuroscience* 18.1, pp. 167–175.

Moliadze, V., A. Antal, and W. Paulus (Dec. 2010). "Electrode-distance dependent after-effects of transcranial direct and random noise stimulation with extra-

cephalic reference electrodes". In: *Clinical Neurophysiology* 121.12, pp. 2165–2171.

Moliadze, V., L. Sierau, et al. (2019). "After-effects of 10Hz tACS over the prefrontal cortex on phonological word decisions". In: *Brain Stimulation* 12.6, pp. 1464–1474.

Moran, R. J. et al. (2010). "Peak frequency in the theta and alpha bands correlates with human working memory capacity". In: *Frontiers in Human Neuroscience* 4.November, pp. 1–12.

Nakazono, H. et al. (2016). "Phase and Frequency-Dependent Effects of Transcranial Alternating Current Stimulation on Motor Cortical Excitability." In: *PloS one* 11.9, e0162521.

Neuling, T., S. Rach, and C. S. Herrmann (2013). "Orchestrating neuronal networks: sustained after-effects of transcranial alternating current stimulation depend upon brain states." In: *Frontiers in human neuroscience* 7, p. 161.

Nguyen, V. et al. (2017). "Regret for expected improvement over the best-observed value and stopping condition". In: *Journal of Machine Learning Research* 77, pp. 279–294.

Nieder, A. (2016). "The neuronal code for number". In: *Nature Reviews Neuroscience* 17.6, pp. 366–382.

Nitsche, M. A. and W. Paulus (Nov. 2001). "Sustained excitability elevations induced by transcranial DC motor cortex stimulation in humans." eng. In: *Neurology* 57.10, pp. 1899–1901.

Ouyang, L. et al. (2020). "Bayesian hierarchical modelling for process optimisation". In: *International Journal of Production Research*.

Ozen, S. et al. (2010). "Transcranial Electric Stimulation Entrains Cortical Neuronal Populations in Rats". In: *Journal of Neuroscience* 30.34, pp. 11476–11485.

Park, J. (2020). "Contextual Bayesian optimization with trust region (CBOTR) and its application to cooperative wind farm control in region 2". In: *Sustainable Energy Technologies and Assessments* 38.February, p. 100679.

Parsons, S. and J. Bynner (2005). "Does Numeracy Matter More ?" In: *National Research and Development Centre for Adult Literacy and Numeracy*, pp. 1–37.

Pasqualotto, A. (2016). "Transcranial random noise stimulation benefits arithmetic skills". In: *Neurobiology of Learning and Memory* May.

Pellegrini, M., M. Zoghi, and S. Jaberzadeh (2018). "Cluster analysis and subgrouping to investigate inter-individual variability to non-invasive brain stimulation: A systematic review". In: *Reviews in the Neurosciences* 29.6, pp. 675–697.

Polanía, R., M. A. Nitsche, and C. C. Ruff (2018). "Studying and modifying brain function with non-invasive brain stimulation". In: *Nature Neuroscience* 21.2, pp. 174–187.

Popescu, T. et al. (2016). "Transcranial random noise stimulation mitigates increased dif fi culty in an arithmetic learning task". In: *Neuropsychologia* 81, pp. 255–264.

Price, A. R. and R. H. Hamilton (May 2015). "A Re-evaluation of the Cognitive Effects From Single-session Transcranial Direct Current Stimulation". In: *Brain Stimulation* 8.3, pp. 663–665.

Radecke, J.-O. et al. (2019). "Simulating individually targeted transcranial electric stimulation for experimental application". In: *bioRxiv*.

Rasmussen, C. and C. Williams (2006). *Gaussian processes for machine learning*. arXiv: 026218253X.

Ridding, M. C. and U. Ziemann (2010). "Determinants of the induction of cortical plasticity by non-invasive brain stimulation in healthy subjects". In: *Journal of Physiology* 588.13, pp. 2291–2304.

Ritchie, S. J. and T. C. Bates (2013). "Enduring Links From Childhood Mathematics and Reading Achievement to Adult Socioeconomic Status". In: *Psychological Science* 24.7, pp. 1301–1308.

Rodriguez-Larios, J. and K. Alaerts (2019). "Tracking Transient Changes in the Neural Frequency Architecture: Harmonic Relationships between Theta and Alpha Peaks Facilitate Cognitive Performance". In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 39.32, pp. 6291–6298.

Ruit, M. van de and M. J. Grey (2019). "False positives associated with responder/non-responder analyses based on motor evoked potentials". In: *Brain Stimulation* 12.2, pp. 314–318.

Russell, M. J. et al. (2017). "Sex and Electrode Configuration in Transcranial Electrical Stimulation". In: *Frontiers in Psychiatry* 8, pp. 1–9.

Santarnecchi, E. et al. (2016). "Individual Differences and Specificity of Prefrontal Gamma Frequency-tACS on Fluid Intelligence Capabilities". In: *Cortex* 75, pp. 33–43.

Sarkar, A. and R. Cohen Kadosh (2016). "Transcranial Electrical Stimulation and Numerical Cognition". In: *Canadian Journal of Experimental Psychology* 70.1, pp. 41–58.

Schecklmann, M. et al. (Mar. 2017). "P056 Daily transcranial random noise stimulation of the temporal cortex in chronic tinnitus". In: *Clinical Neurophysiology* 128.3, e33–e34.

Schönbrodt, F. D. and A. M. Stefan (2018). *BFDA: An R package for Bayes factor design analysis (version 0.4.0)*.

Schulz, E., E. Konstantinidis, and M. Speekenbrink (2018). "Putting bandits into context: How function learning supports decision making". In: *Journal of Experimental Psychology: Learning Memory and Cognition* 44.6, pp. 927–943.

Sheffield, J. et al. (Jan. 2020). "How can noise alter neurophysiology in order to improve human behaviour? A combined transcranial random noise stimulation and electroencephalography study". In: *bioRxiv*, p. 2020.01.09.900118.

Snoek, J., H. Larochelle, and R. Adams (2012). "Practical Bayesian Optimization of Machine Learning Algorithms." In: *Nips*, pp. 1–9. arXiv: 1206.2944.

Snowball, A. et al. (2013). "Long-term enhancement of brain function and cognition using cognitive training and brain stimulation". In: *Current Biology* 23.11, pp. 987–992.

Strube, W. et al. (2016). "Bidirectional variability in motor cortex excitability modulation following 1 mA transcranial direct current stimulation in healthy participants". In: *Physiological Reports* 4.15, e12884.

Strüber, D. et al. (2015). "On the possible role of stimulation duration for after-effects of transcranial alternating current stimulation". In: *Frontiers in Cellular Neuroscience* 9, p. 311.

Taatgen, N. A. (2013). "The nature and transfer of cognitive skills." In: *Psychological review* 120.3, pp. 439–471.

Tavakoli, A. V. and K. Yun (2017). "Transcranial Alternating Current Stimulation (tACS) Mechanisms and Protocols". In: *Frontiers in cellular neuroscience* 11.September, pp. 1–10.

Teo, J. T. et al. (2014). "Late cortical plasticity in motor and auditory cortex: Role of met-allele in BDNF Val66Met polymorphism". In: *International Journal of Neuropsychopharmacology* 17.5, pp. 705–713.

Terney, D. et al. (2008). "Increasing Human Brain Excitability by Transcranial High-Frequency Random Noise Stimulation". In: *Journal of Neuroscience* 28.52, pp. 14147–14155.

Tervo, A. E. et al. (2020). "Automated search of stimulation targets with closed-loop transcranial magnetic stimulation". In: *NeuroImage* 220.March, p. 117082.

Treasure, T. and K. D. MacRae (1998). "Minimisation: the platinum standard for trials? Randomisation doesn't guarantee similarity of groups; minimisation does." In: *BMJ* 317.7155, pp. 362–363.

Tseng, P. et al. (2012). "Unleashing potential: transcranial direct current stimulation over the right posterior parietal cortex improves change detection in low-performing individuals". In: *The Journal of Neuroscience* 32.31, pp. 10554–61.

Vannorsdall, T. D. et al. (Mar. 2016). "Reproducibility of tDCS Results in a Randomized Trial: Failure to Replicate Findings of tDCS-Induced Enhancement of Verbal Fluency". In: *Cognitive And Behavioral Neurology* 29.1, pp. 11–17.

Vöröslakos, M. et al. (2018). "Direct effects of transcranial electric stimulation on brain circuits in rats and humans". In: *Nat Commun* in Press.

Vosskuhl, J., R. J. Huster, and C. S. Herrmann (2016). "BOLD signal effects of transcranial alternating current stimulation (tACS) in the alpha range: A concurrent tACS–fMRI study". In: *NeuroImage* 140, pp. 118–125.

Wagenmakers, E. J., H. L. Van Der Maas, and R. P. Grasman (2007). "An EZ-diffusion model for response time and accuracy". In: *Psychonomic Bulletin and Review* 14.1, pp. 3–22.

Wiethoff, S., M. Hamada, and J. C. Rothwell (May 2014). "Variability in Response to Transcranial Direct Current Stimulation of the Motor Cortex". In: *Brain Stimulation* 7.3, pp. 468–475.

Willmes, K. (2008). "Chapter 17 Acalculia". In: *Handbook of Clinical Neurology* 88, pp. 339–358.

Wörsching, J. et al. (2018). "Testing assumptions on prefrontal transcranial direct current stimulation: Comparison of electrode montages using multimodal fMRI". In: *Brain Stimulation*.

Appendix A

A.1 Screening and consent form

University of Oxford
Department of Experimental Psychology
South Parks Road
Oxford OX1 3UD

CONFIDENTIAL

Participant Consent Form

Project Title: Modulating brain activity in psychological tasks

• Do you currently have any of the following fitted to your body? ☐ Heart pacemaker ☐ Medication pump ☐ Any other metallic object ☐ Cochlear implant ☐ Surgical clips	YES	NO
• Do you have a personal or family history of epileptic fits or seizures? (If **YES**, then please state your relationship to the affected family member.)	YES	NO
• Do you have a personal history of migraines?	YES	NO
• Have you ever suffered from febrile convulsions in infancy, or fainted?	YES	NO
• Have you ever had a head trauma that was diagnosed as a concussion or was associated with a loss of consciousness? (If **YES**, please describe.)	YES	NO
• Do you have any other medical/psychiatric/neurological disorder?	YES	NO
• Do you have any metal in your brain or elsewhere in your body? (If **YES**, please specify item and location.)	YES	NO
• Have you ever undergone a neurosurgical procedure (including eye surgery)? (If **YES**, please describe.)	YES	NO
• Could you be pregnant?	YES	NO
• Are you currently taking any prescribed or unprescribed medication (or herbal remedies), or participated in the last month in a pharmacological experiment? (If **YES**, please list.)	YES	NO
• Have you drunk any alcohol or used any recreational drugs in the last 24 hours?	YES	NO
• Have you had more than one cup of coffee, or other sources of caffeine, in the last 6 hours?	YES	NO
• When did you have your last proper meal?		
• How many cups of fluids (water/juice/milk/soft drinks etc.) have you had today?		
• Have you removed all metallic items (e.g. jewellery, coins, watches) from yourself?	YES	NO
• Have you had a full night's sleep (≥6 hours) the night before the study?	YES	NO
• Have you participated in a TMS study in the last 7 days?	YES	NO
• Have you participated in a tDCS study in the last 7 days? (If **YES**, when was that: ____ Days)	YES	NO
• Do you understand that you should not drive for one hour after receiving tDCS/TMS?	YES	NO
• Are you left or right handed?	Left	Right
• Date of Birth	/ /	

Please initial each box

I confirm that I have read and understand the information sheet for the above study. I have had the opportunity to consider the information, ask questions and have had these answered satisfactorily.	☐
I understand that my participation is voluntary and that I am free to withdraw at any time, without giving any reason, without my medical care or legal rights being affected.	☐
I understand that relevant sections of data collected during the study may be looked by individuals, from regulatory authorities, where it is relevant to my taking part in this research. I give permission for these individuals to have access to my records.	☐
I understand that relevant data collected during the study may be looked at by individuals from the University of Oxford for the purpose of audit and monitoring, and where it is relevant to my taking part in this research. I give permission for these individuals to have access to my records.	☐
I understand who will have access to personal data provided, how the data will be stored and what will happen to the data at the end of the project.	☐
I understand how to raise a concern or make a complaint.	☐
I understand that this project has been reviewed by, and received ethics clearance through, the University of Oxford Central University Research Ethics Committee.	☐
I agree for my personal data to be kept in a secure database for the purpose of contacting me about future studies.	☐
I agree to take part in the above study.	☐

If you are interested in receiving a general feedback sheet about the results of the study please inform the experimenter.
This study has been reviewed and given favourable opinion by the IDREC-C2-2014-033.

Signed_____ Date_____

Name of volunteer (block letters)_____

Name of Investigator_____

Signature of Investigator_____

Signature of Senior Investigator***** _____

***** Required for **one** or more highlighted responses.

A.2 Sensation questionnaire

The following questionnaire was given to participants to determine the sensations that they experienced in response to tES and was adapted from Fertonani, Rosini, et al. (2010). The first question was answered in response to each of the following sensations; Itching, Pain, Burning, Warmth/Heat, Fatigue/Decreased alertness, Flashing lights (other than the stimulus). All other questions were only answered once each time the questionnaire was presented.

1. Please indicate whether you experienced any discomfort during the stimulation by typing the corresponding number on a scale from 1 (no pain) to 10 (worst pain imaginable):
    - 1 = None (I did not feel the sensation)
    - 2–3 = Mild (I mildly felt the sensation)
    - 4–6 = Moderate (I felt the sensation)
    - 7–10 = Severe (I felt the sensation to a considerable degree)
2. In the case of perceived sensations: How much did these sensations affect your general state?
    - 1 = Not at all
    - 2 = Slightly
    - 3 = Considerably
    - 4 = Much
    - 5 = Very much
3. Do you believe that you received real or placebo stimulation?
    - 1 = real
    - 2 = placebo
    - 3 = I do not know