

Sensorimotor adaptation: mechanisms, modulation and rehabilitation potential



Pierre Petitet
New College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary Term 2018

Sensorimotor adaptation: mechanisms, modulation and rehabilitation potential

Pierre Petitet

Hilary Term 2018

New College

*A thesis submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy at the University of Oxford*

Abstract

Adaptation is a fundamental property of the nervous system that underlies the maintenance of successful actions through flexible reconfiguration of sensorimotor processing. The primary aims of this thesis are 1) to investigate the computational and neural underpinnings of sensorimotor memory formation during prism adaptation (PA) in humans, and 2) how they interact with anodal transcranial direct current stimulation (a-tDCS) of the primary motor cortex (M1), in order to 3) improve efficacy of prism therapy for post-stroke spatial neglect.

In [chapter 4](#), we modify an influential state-space model of adaptation in order to characterise the contribution of short and long memory timescales to motor behaviour as sensorimotor after-effects (AEs) develop during PA. This enables us, in the multimodal 7 Tesla MRI experiment reported in [chapter 5](#), to demonstrate that the level of M1 excitation:inhibition causally sets the relative contribution of long versus short memory timescales during PA, thus determining behavioural persistence of the AE at retention in young healthy adults. This finding offers a bridge between different levels of investigation by providing a biologically plausible neuro-computational model of how sensorimotor memories are formed and enhanced by a-tDCS. In [chapter 6](#), we use the ageing motor system as a model of reduced GABAergic inhibition and show that the age-related decrease in M1 GABA explains why older adults demonstrate more persistent prism AEs. Taken together, these data indicate that the reduction in M1 GABAergic inhibition via excitatory a-tDCS during PA has the potential to enhance persistence of adaptation memory in both young and older adults.

Informed by these results, we subsequently ask whether standard (multi-session) PA therapy combined with left M1 a-tDCS translates to greater and/or longer-lasting clinical improvements in post-stroke spatial neglect patients. In [chapter 7](#), we compare the multimodal neuroimaging data of six neglect patients to normative data of age-matched controls. We show that in all patients, the lesion interrupted long-range frontoparietal connections, and we provide direct evidence for a pathological left dominance of activity within the lateral occipital cortex during deployment of bilateral visuospatial attention. In [chapter 8](#), we present the behavioural performance of these patients throughout the two phases of the clinical study (i.e. before and after either PA + real M1 a-tDCS or PA + sham M1 a-tDCS). There was no clear effect of a-tDCS on the therapeutic effect of PA in these patients.

The results of the studies presented in this thesis provide a novel insight into the neuro-computational mechanisms of sensorimotor memory formation and its modulation by a-tDCS in the healthy brain. Further investigation of how these mechanisms relate to therapeutic improvements following PA in certain neglect patients is needed.

This thesis contains approximately 50,000 words.

Acknowledgements

First and foremost, I would like to express my sincere gratitude and thanks to a number of people who have made my time in Oxford so unique and unforgettable in many different ways:

To the participants, for enduring long hours of testing and scanning, and for having the patience and stamina to keep coming back to see the studies through. In particular to the patients and their incredible motivation – working with them was a constant source of astonishment and reminder of the modesty yet importance of our research.

To my truly amazing co-supervisors Jacinta O’Shea and Heidi Johansen-Berg, for giving me the life-changing opportunity to carry out my research project in one of the best neuroscience laboratories one can wish to work in. I feel honoured and fortunate to have worked with them and this thesis would not have been without their generous support. To Heidi in particular for supervising me on her own during my second year, and for being always so enthusiastic and resourceful. To Jacinta in particular for interesting me enough in her work to put on hold my medical studies and focus entirely on research for a few years, and for pushing me to develop my own line of thinking throughout my DPhil.

A special mention goes to Jill O’Reilly, for guiding me in my exploration of computational neurosciences – from teaching me the fundamentals of Matlab and FSL to brainstorming about state-space and Bayesian modelling.

To the following collaborators: MaryAnn Noonan, who acquired the data presented in [Appendix B](#), introduced me to UNIX and showed me how to make pretty figures; Holly Bridge, for insightful feedback on the manuscript reported in [Appendix B](#); Shigeru Kitazawa sensei, who generously shared his knowledge on state-space modelling, provided helpful comments on the work reported in [chapter 5](#) and invited me to speak at the 1st International Symposium on Systems Intelligence Division in Osaka in January 2018; Uzay Emir, who provided assistance with the magnetic resonance spectroscopy analysis presented in [chapter 5](#); Damien Pastor, who collected the prism adaptation data of the first experiment reported in [chapter 4](#); David Henderson-Slater and Udo Kischka, who kindly helped with patient recruitment.

To Ana Margarida Silva Goncalves, for her assistance collecting data for the “GABAptation” project ([chapter 5](#)), and most importantly for teaching me how to be a supervisor and a friend at the same time.

To Laurie Josephs and Tom Smekja, for their assistance in the “PriStim” study ([chapter 7](#) & [8](#)). To Laurie in particular, for being so patient and supportive with me when I was a fresher in the laboratory and spoke *Franglish* only.

To Rebecca Annells and Gershon Spitz, for their assistance collecting data for the “PAgeing” study ([chapter 6](#)). To Gershon in particular, for showing so much interest in my work and for offering to keep some projects running when I will leave the laboratory.

To Piergiorgio Salvan, for his help with the resting-state analysis in [chapter 5](#), and for the countless enthusiastic complast (computational plasticity) chats we have had over the last year and will have in the future.

To Chris Gallagher, for his help building the automated prism adaptation set-up. He gave me countless advises on touch screens, liquid crystal shutters and Arduinos, wrote a program for the big red button, and even helped weld bits and pieces together.

To the Marie Skłodowska-Curie actions (European Commission), for providing my funding. It was also a truly unique experience to meet regularly with other fellows from the Adaptive Brain Computation (ABC) network. In particular, many thanks to Heather, Sinead, Ling, Poly, Nuno and Joey for making these meetings so entertaining and memorable.

To Harriet Dempsey-Jones, for proof-reading my thesis and for being such an amazing matey in general. Although we both came “for six months”, it looks the time is ripe for us to leave now.

To Naiara Demnitz and Sanne Kikkert, for being the best DPhil buddies one could ask for. I will miss you guys!

To all members of the Plasticity group, Neuroplastics and everyone at FMRIB, for being who they are: a fantastic group of people to be around and work with. I am thankful to have found in this laboratory the perfect balance of excellence and fun that is, in my opinion, the best formula for success.

To all alumni and staff of New College, for creating such a wonderful environment to be part of, and for bringing a bit of magic to pretty much anything they organise.

To Yves Rossetti for encouraging me to pursue a MD-PhD training, advising me ever since, and being a wonderful source of inspiration.

To the *École de l’Inserm Liliane Bettencourt*, for facilitating and accompanying me in my MD-PhD training since 2012.

To Hiyama sensei and all members of the Oxford Kendo club, for helping me let-

ting off some steam with some kirikaeshi, suburi and jigeiko.

To my *Auberge espagnole* – William Parker, Maria Reyes, Chiti (Bang Bang) Garcia Ayuso, Pepe Silvera Ramon, Ari Vazquez Portero, Hugo Calderon, Anica Maan, and Claudette! Many thanks for having been brave enough to welcome me in the *102 Hurst street* crew when I was only 21 years old, and for having been part of my family ever since.

To my parents, Sylvain and Véronique, and my sister, Marie, for their endless support all along the way, and for being great role models.

Finally, gracias, gràcies, merci and thanks to Mire, for your unconditional love and support, for crossing two seas to make Oxford our home, for the joy, carnaval, fishes, chistorra, fireworks, spices, waves, bubbles, flamenco you bring to my life. I hope that when we will be old and wise, and this thesis will be sitting on a dusty shelf, we will remember our chapter in Oxford as the preamble to many more adventures in Madagascar, Montpellier, Barbate, and many other places I cannot even imagine yet. Te amo.

This thesis is dedicated to Mire,
with all my love

Publications arising from this thesis and supplementary work

I have dedicated half of the first year of my doctoral studies to analysing imaging and behavioural data from a combined transcranial magnetic stimulation/functional magnetic resonance imaging dataset previously collected in humans by Dr MaryAnn Noonan. Working on this project was very useful in terms of developing my image analysis skills and I had enough intellectual freedom to provide a substantial contribution to the study. The manuscript arising from this work was published in the journal *Neuropsychologia* in July 2015. The objective of this study (finding behavioural and neuro-physiological evidence of the inter-hemispheric competition account of visual extinction in the healthy brain) is in keeping with the general theme of this thesis but is not directly related to the central working hypothesis. This work has therefore been omitted from the main text of this thesis but is provided in [Appendix B](#) and in published form under the following citation:

Petit, P., Noonan, M. P., Bridge, H., O'Reilly, J. X. & O'Shea, J. Testing the inter-hemispheric competition account of visual extinction with combined TMS/fMRI. *Neuropsychologia* **74**, 63–73 (2015)

In preparation of this thesis, I, together with my supervisors, have prioritised the development of manuscripts for submission. As such, the general introduction ([chapter 1](#)) and experimental chapters [5](#) & [6](#) are adapted from manuscripts that I previously prepared or am in the process of preparing for publication in peer-reviewed scientific journals. The respective citations are detailed at the start of each chapter.

Table of Contents

Abstract	i
Acknowledgements	ii
Publications arising from this thesis and supplementary work	vi
List of Figures	xv
List of Tables	xix
List of Acronyms	xxii
1 Introduction	1
1.1 Why study prism adaptation?	2
1.2 Prism adaptation procedures	5
1.3 Rehabilitation potential for spatial neglect	8
1.4 The traditional dual-process framework: strategic control versus spatial realignment	12
1.4.1 Theoretical framework	12
1.4.2 Behavioural predictions	14
1.4.3 Strategic control	15
1.4.4 Spatial realignment	19
1.4.5 Why is this framework not satisfying?	21
1.5 An alternative theoretical framework: multiple timescale models . . .	23
1.6 Conclusion	27
1.7 Aims of this thesis	28
2 Methods for stimulating the brain	30
2.1 Transcranial Magnetic Stimulation (TMS)	30
2.1.1 Practical applications	31

2.1.2	Motor Evoked Potentials and resting motor threshold	32
2.1.3	Repetitive TMS	33
2.1.4	Safety considerations	35
2.2	Transcranial Direct Current Stimulation (tDCS)	38
2.2.1	Practical applications	39
2.2.2	Physiological basis of tDCS	41
2.2.2.1	Electrophysiological effect of Direct Current stimulation (DCS) <i>in vitro</i>	42
2.2.2.2	Mapping the current inside the head	45
2.2.2.3	Electrophysiological effect of tDCS <i>in vivo</i>	46
2.2.3	Effect of M1 anodal tDCS on sensorimotor adaptation	49
2.2.4	Safety and ethical considerations	51
3	Methods for imaging the brain	53
3.1	Physics principles of Magnetic Resonance Imaging	53
3.1.1	Polarisation and precession	54
3.1.2	RF pulses and excitation	55
3.1.3	Relaxation and signal generation	56
3.1.4	Read-out methods	58
3.1.5	Pulse sequences and image contrast	59
3.2	Diffusion-weighted Imaging	61
3.2.1	What is diffusion?	61
3.2.2	Diffusion-weighted MRI sequences	62
3.2.3	Image artefacts: distortions and eddy current	63
3.2.4	Diffusion Tensor Imaging and derived scalars	64
3.3	Functional Magnetic Resonance Imaging	66
3.3.1	Haemodynamic response and blood oxygen-level dependent (BOLD) signal	66
3.3.2	fMRI acquisition and artefacts	67
3.3.3	Statistical analysis of task fMRI data	68
3.3.4	Statistical analysis of resting state fMRI data	71
3.4	Magnetic Resonance Spectroscopy	74
3.4.1	Principles of <i>in vivo</i> MR spectroscopy	74
3.4.2	Acquisition considerations: water suppression, outer volume suppression and shimming	75
3.4.3	Standard spectrum analysis and metabolite quantification	76

3.4.4	Quality filtering	78
3.4.4.1	Comment on the use of the relative Cramér-Rao Lower Bounds	78
3.4.4.2	Is there a better alternative?	80
4	Inferring computational states driving adaptation dynamics	82
4.1	Introduction	83
4.2	Material and Methods	86
4.2.1	Prism adaptation protocol	86
4.2.2	Modelling framework	90
4.2.2.1	Existing state-space models of adaptation	90
4.2.2.2	Weighted two-state model	93
4.2.3	Data acquisition technique	95
4.2.4	Experimental procedures	96
4.3	Results	98
4.3.1	Model comparison during the adaptation phase	98
4.3.2	Predictive power during active de-adaptation	102
4.3.3	Predictive power 24 hours after the adaptation	104
4.4	Discussion	106
5	Causal explanation of individual differences in human sensorimotor memory formation	109
5.1	Introduction	109
5.2	Material and Methods	112
5.2.1	Participants	112
5.2.2	Experimental design	112
5.2.3	Prism adaptation procedure	113
5.2.4	Transcranial direct current stimulation	116
5.2.4.1	During behavioural sessions	116
5.2.4.2	Inside the MRI scanner	117
5.2.5	Statistical analysis of raw prism adaptation data	118
5.2.6	Computational analysis of prism adaptation data	119
5.2.7	Regression analyses	122
5.2.8	BDNF Val66Met polymorphism	122
5.2.9	MR data acquisition	122
5.2.9.1	T1-weighted structural images	123
5.2.9.2	MR spectroscopy	124

5.2.9.3	Diffusion weighted images (DWI)	124
5.2.9.4	Resting state functional MRI (rs-fMRI)	125
5.2.10	MR data analysis	125
5.2.10.1	MR spectroscopy	125
5.2.10.2	Diffusion-weighted images	128
5.2.10.3	Resting state functional MRI	130
5.3	Results	132
5.3.1	Computational model of sensorimotor memory formation	132
5.3.2	Functional and structural connectivity correlates of memory persistence	135
5.3.3	Sensorimotor cortex excitation:inhibition ratio causally determines memory persistence	137
5.3.4	Enhanced memory persistence via sensorimotor cortex stimulation is modulated by genotype	141
5.3.5	Supplementary analyses (group mean effects)	143
5.3.5.1	Statistical analysis of motor behaviour during and after prism adaptation	143
5.3.5.2	Statistical analysis of the neurochemical effect of left M1 a-tDCS	145
5.3.5.3	Statistical analysis of the a-tDCS effect on resting-state networks	146
5.4	Discussion	147
5.5	Conclusions	151
6	Sensorimotor memory formation in the ageing human brain	152
6.1	Introduction	152
6.2	Material and methods	155
6.2.1	Participants	155
6.2.2	Experimental procedures	157
6.2.2.1	Baseline prism adaptation experiment	157
6.2.2.2	MR spectroscopy experiment	158
6.2.2.3	Brain stimulation experiment	159
6.2.3	Transcranial direct current stimulation (tDCS)	159
6.2.4	MRS data analysis	160
6.2.5	Computational analysis of prism adaptation data	163
6.2.6	Statistical analysis	165

6.3	Results	166
6.3.1	Comparison of prism adaptation data between young and older adults	166
6.3.2	Relationship between excitation:inhibition and relative weight in older adults	171
6.3.3	Older adults show more persistent after-effects	172
6.3.4	Older adults have lower GABAergic inhibition	173
6.3.5	Lower M1 GABA is associated with more persistent after-effect	175
6.3.6	Mediation analysis	176
6.3.7	Effect of online M1 a-tDCS on PA in elderly adults	179
6.3.8	Neurotransmitter levels in neglect	181
6.4	Discussion	182
6.5	Conclusions	185
7	Functional and structural basis of spatial neglect	186
7.1	Introduction	186
7.2	Material and methods	190
7.2.1	Participants	190
7.2.2	Posner-like task	191
7.2.3	MR data acquisition	192
7.2.3.1	Resting state fMRI	193
7.2.3.2	Posner-like task fMRI	193
7.2.3.3	T1-weighted structural images	193
7.2.3.4	Diffusion weighted imaging	194
7.2.4	MR data analysis	194
7.2.4.1	Lesion mapping	194
7.2.4.2	Posner-like task fMRI data	194
7.2.4.3	Resting state fMRI data	196
7.2.4.4	Diffusion-weighted data	197
7.3	Results	198
7.3.1	Lesion mapping	198
7.3.2	Re-organisation of the left frontoparietal network in spatial neglect	201
7.3.3	Behavioural performance on the Posner-like task	202
7.3.4	Altered inter-hemispheric balance in spatial neglect	205
7.3.4.1	Age-matched healthy individuals	205

7.3.4.2	Neglect patients	209
7.3.4.3	Neglect patients versus age-matched healthy controls	209
7.4	Discussion	213
7.5	Conclusions	218
8	Combining PA and tDCS to improve neglect rehabilitation	220
8.1	Introduction	220
8.2	Material and Methods	223
8.2.1	Participants and recruitment process	223
8.2.2	Experimental design	225
8.2.3	Assessment of spatial neglect	226
8.2.3.1	Neglect test battery	227
8.2.3.2	Posner cueing task	229
8.2.4	Intervention procedure	231
8.2.5	Transcranial direct current stimulation	232
8.2.6	Neglect test scoring	233
8.2.6.1	Neglect score	233
8.2.6.2	Posner cueing task	235
8.2.7	Analysis of performance during the intervention	236
8.2.8	Computational modelling of the intervention	237
8.2.9	Statistical analysis	239
8.3	Results	239
8.3.1	Effect of the intervention on clinical measures of spatial neglect	239
8.3.1.1	Primary outcome: neglect score	240
8.3.1.2	Secondary outcome: Posner cueing task	242
8.3.2	Patients adapted to prisms and demonstrated a sensorimotor after-effect	244
8.3.3	Effect of left M1 anodal tDCS on prism adaptation in patients	245
8.3.4	Extracting hidden computational states during the intervention	248
8.4	Discussion	251
8.5	Conclusions	253
9	Conclusions and future directions	254
9.1	Aims of the thesis	254
9.2	Summary of key findings	255
9.3	Limitations and future directions	260

9.3.1	Does M1 a-tDCS enhances clinical efficacy of prism adaptation for neglect rehabilitation?	260
9.3.2	Towards a model of prism adaptation that spans Marr's three levels of analysis	261
9.3.3	A better characterisation of the fast and slow systems is needed	263
9.3.4	Explaining inter-individual differences in sensitivity to brain stimulation	265
9.4	General conclusions	266
A	Bayesian Decision Theory	267
A.1	Internal models for sensorimotor control	267
A.2	Decision theory and rational movement planning	269
A.3	Bayesian statistics and optimal state estimation	275
A.4	Internal model updating and credit assignment	280
B	Testing the inter-hemispheric competition account of visuospatial attention with combined TMS/fMRI	286
B.1	Introduction	287
B.2	Material and methods	290
B.2.1	Participants	290
B.2.2	TMS functional localisation of right posterior parietal cortex .	291
B.2.3	Transcranial Magnetic Stimulation	294
B.2.4	Visual Extinction task	295
B.2.5	Experimental Procedures	297
B.2.5.1	Experiment 1: Behavioural effect of right parietal TMS on extinction task	297
B.2.5.2	Experiment 2: TMS/fMRI of the extinction task . .	298
B.2.6	MR image acquisition	298
B.2.7	Functional magnetic resonance imaging data analysis	299
B.2.8	Statistical analysis of behavioural data	300
B.3	Results	302
B.3.1	Behavioural data	302
B.3.1.1	Experiment 1: Behavioural effect of right parietal TMS on the extinction task	302
B.3.1.2	Experiment 2: TMS/fMRI of the extinction task . .	303
B.3.2	Functional brain imaging data	305
B.3.2.1	Extinction task network	305

B.3.2.2	Effect of TMS on the inter-hemispheric balance . . .	307
B.3.2.3	Inter-individual variation in the magnitude of the TMS effect	308
B.4	Discussion	312
B.5	Conclusions	318
C	Supplementary information from Chapter 6	319
	Bibliography	322

List of Figures

1.1	Prism adaptation	7
1.2	Traditional dual-process framework	13
1.3	The three timescales state-space model for prism adaptation	25
2.1	Example TMS machine and application	31
2.2	Effect of repetitive TMS on MEPs	34
2.3	Example tDCS kit and application	38
2.4	Left M1 anodal tDCS montage	40
2.5	DCS promotes LTP in motor cortical slices	43
2.6	DCS-LTP requires NMDA and BDNF	44
2.7	Simulation of the spatial distribution of tDCS-induced intracranial electric field	45
2.8	Effect of tDCS on cortical excitability in humans	47
2.9	Change in GABA before, during, and after anodal tDCS	48
3.1	Magnetic moments of protons	55
3.2	Free induction decay	57
3.3	Principles of spatial encoding	58
3.4	Example pulse Spin-Echo sequence	60
3.5	Water diffusion in brain tissue	61
3.6	Diffusion-weighted spin echo sequence	62
3.7	Diffusion tensor	65
3.8	BOLD signal timecourse and optimal TE	68
3.9	General Linear Model (GLM)	70
3.10	Independent Component Analysis (ICA)	72
3.11	Resting State Networks (RSNs)	73
3.12	Components of a spectrum	77
3.13	Relationship between concentration estimate and relative Cramér-Rao lower bounds	79

4.1	Prism adaptation protocol	87
4.2	Automated prism adaptation set-up	95
4.3	Experimental protocols of experiment 1	97
4.4	Fit of the two models across the four datasets	100
4.5	Comparison of the Akaike's Information Criterion among the two models for individual fits	102
4.6	Descriptive and predictive value of the weighted two-state model . . .	103
4.7	After-effect memory trace 24 hours after prism adaptation (Experiment 1)	105
4.8	The dynamics of long-term memory reactivation reflect the relative weighting between fast and slow systems	106
5.1	Experimental set-up and protocols	113
5.2	Magnetic Resonance Spectroscopy data	127
5.3	Thirteen resting-state networks extracted from the 25-component analysis	131
5.4	Computational modelling of motor behaviour during prism adaptation	134
5.5	Structural and functional connectivity correlates of memory persistence	136
5.6	Predictions and correlations of MRS, computational and retention measures at baseline	138
5.7	Predictions and correlations of stimulation-induced change in MRS, computational and retention measures	140
5.8	The BDNF Val66Met polymorphism influences the effect of anodal tDCS on retention	142
5.9	Behavioural effect of left M1 a-tDCS	144
5.10	No group mean effect of stimulation on metabolite concentrations . .	146
6.1	Magnetic Resonance Spectroscopy data quality	162
6.2	Prism adaptation and retention data	167
6.3	Model fit in young and older participants	169
6.4	Relationship between M1 excitation:inhibition and relative weight . .	172
6.5	Effect of age AE persistence	173
6.6	Effect of age on spectroscopy measurements	174
6.7	Effect of M1 neurotransmitters on AE persistence	175
6.8	Effect of M1 a-tDCS during prism adaptation	180
6.9	Neurotransmitter levels in neglect	182

7.1	MRI scanning protocol	193
7.2	Patients' lesion anatomy	199
7.3	Functional re-organisation of the left frontoparietal network	201
7.4	Behavioural performance inside the scanner	204
7.5	Activation during the Posner-like task	206
7.6	Motor-related activation in age-matched healthy individuals	207
7.7	Functional lateralisation during the Posner-like task	208
7.8	Comparison of brain activity during the Posner-like task	211
7.9	Results of a tractography run from the region showing increased left dominance in neglect patients	212
8.1	Recruitment flowchart	223
8.2	Experimental design	226
8.3	Example pen-and-paper tests	228
8.4	Posner cueing task	230
8.5	Centre of performance	234
8.6	Schematic representation of the online, offline and total effects	236
8.7	Group mean and individual neglect scores	241
8.8	Behavioural performance on the Posner cueing task	243
8.9	Prism adaptation profile	244
8.10	Time course of pointing accuracy throughout the study	245
8.11	Online (within-day) and offline (between-day) effects of a-tDCS	247
8.12	Fit of the multi-day weighted two-state model to the intervention PA data	249
A.1	Decision theory	272
A.2	Bayesian statistics in sensorimotor control	277
A.3	A Bayesian account of prism adaptation	283
B.1	Functional localisation of parietal cortex and the visual extinction task	293
B.2	Behavioural results of Experiment 1	304
B.3	Cortical network activated during bilateral attention (Experiment 2)	306
B.4	Effect of rTMS on inter-hemispheric balance during bilateral attention	309
B.5	Effect of TMS on brain activity in the region showing a significant leftward shift in inter-hemispheric balance	310
B.6	Baseline-dependency of the effect of TMS on parietal inter-hemispheric balance	311

B.7 Timecourse of TMS-induced change in inter-hemispheric balance . . .	315
---	-----

List of Tables

1.1	Overview of the lesion studies of prism adaptation	18
1.2	Neuroimaging correlates of strategic control	19
1.3	Neuroimaging correlates of spatial realignment	20
2.1	Summary of studies investigating the effect of primary motor cortex anodal tDCS during sensorimotor adaptation.	50
4.1	Model comparison for the four datasets	101
5.1	Diagram of data completeness	123
5.2	MRS characteristics	126
6.1	Participation in each of the three experiments	156
6.2	MRS characteristics	161
6.3	M1 GABA mediates the relationship between age and long-term sen- sorimotor memory	178
7.1	Patients' lesion anatomy	200
7.2	Brain regions showing differences in activity between neglect patients and age-matched healthy individuals	210
8.1	Patient characteristics	225
B.1	Network of brain regions significantly activated by the extinction task	307

Acronyms

ACC Anterior Cingulate Cortex

AE after-effect

AIC Akaike's Information Criterion

ANG Angular gyrus

ARAT Action Research Arm Test

ATP Adenosine triphosphate

BDNF Brain-derived Neurotrophic Factor

BOLD Blood Oxygen-level Dependent

CBS Catherine Bergego Scale

CLP Closed Loop Pointing

CoP Centre of Performance

CSF Cerebrospinal Fluid

CST Corticospinal tract

DAN Dorsal Attention Network

DC Direct Current

DTI Diffusion Tensor Imaging

DWI Diffusion-weighted Imaging

E:I excitation:inhibition

EMG electromyographic

EPI Echo-planar Imaging

EPSP excitatory postsynaptic potential

FA	Fractional Anisotropy
FEF	Frontal Eye Field
fEPSP	field excitatory postsynaptic potentials
FG	Fusiform Gyrus
fMRI	functional Magnetic Resonance Imaging
FP	Frontal Pole
FPN	Frontoparietal network
FSL	FMRIB software library
GABA	gamma-Aminobutyric acid
GLM	General Linear Model
Gln	Glutamine
Glu	Glutamate
Glx	Glutamix (Glu + Gln)
GM	Grey Matter
ICA	Independent Component Analysis
ICC	Intraclass Correlation Coefficient
IFG	Inferior Frontal Gyrus
IFOF	Inferior Fronto-occipital Fasciculus
Ins	Insula
IPS	Intra-parietal sulcus
LFS	Low Frequency Stimulation
LI	Lateralisation index
LOC	Lateral Occipital Sulcus
LTD	Long-Term Depression
LTP	Long-Term Potentiation
M1	Primary motor cortex
MD	Mean Diffusivity

MEP	motor-evoked potential
MFG	Middle Frontal Gyrus
MNI	Montreal Neurological Institute
MRI	Magnetic Resonance Imaging
MRS	Magnetic Resonance Spectroscopy
MTG	Middle Temporal Gyrus
NMDA	N-methyl-D-aspartate
NTB	Neglect Test Battery
OLP	Open Loop Pointing
PA	prism adaptation
RCT	Randomised Clinical Trial
rMT	resting motor threshold
RSN	Resting State Network
RT	Reaction Time
S1	Primary sensory cortex
SLF	Superior Longitudinal Fasciculus
SMA	Supplementary Motor Area
SPL	Superior Parietal Lobe
STS	Superior Temporal Sulcus
rCr	total Creatine
tDCS	Transcranial direct current stimulation
TMS	Transcranial Magnetic Stimulation
TPJ	Temporoparietal junction
V1	Primary visual cortex
VAN	Ventral Attention Network
WM	White Matter

Chapter 1

Introduction

The generation of an appropriate movement is arguably the end product of most brain operations. Our interaction with the world (e.g. reading, talking, eating, texting, etc.) fundamentally relies on our ability to generate tightly controlled actions. In an ideal (deterministic) world, the same command would always produce the exact same movement, so the problem of motor control would be an easy one to solve. However in reality, many disturbances can alter the relationship between a desired goal and the motor command required to achieve it. This observation underpins the need for the brain to be equipped with a mechanism that allows the flexible reconfiguration of sensorimotor processing to exert motor control that is robust to disturbances. This is the central theme of this thesis. This chapter provides an introduction to this research topic, with a particular emphasis on the computations involved and the potential applications for post-stroke rehabilitation of spatial neglect. The core of this chapter is adapted from:

Petitot, P., O'Reilly, J. X. & O'Shea, J. Towards a neuro-computational account of prism adaptation. *Neuropsychologia* (2017)

1.1 Why study prism adaptation?

Adaptation is a fundamental property of the nervous system that enables organisms to flexibly reconfigure sensorimotor processing to counteract perturbations that cause performance errors (Shadmehr *et al.* 2010; Franklin & Wolpert 2011). Consider, for example, the case of a basketball player shooting at various times throughout a game. As the game progresses, so muscles will fatigue, such that the same motor command produces a different outcome from one shoot to another. A lateral wind might also alter the trajectory of the ball and deviate it from the aimed basket. In these two situations, an internal (muscle fatigue) or external (wind) disturbance introduces systematic deviations from the intended action goal. These perturbations require the relationship between a desired action goal and the motor commands that execute it to be reconfigured, to avoid the large systematic errors in performance that would ensue if the nervous system were unable to adapt and correct for the perturbations. Thus, adaptation underwrites the maintenance of successful actions across the lifespan.

In a laboratory context, sensorimotor adaptation has been studied experimentally using a variety of methods (e.g. visuomotor rotation, force-field adaptation, saccade adaptation, Coriolis forces, etc.) (Lackner & Dizio 1994; Shadmehr & Mussa-Ivaldi 1994; Mazzoni & Krakauer 2006; Ethier *et al.* 2008a). This thesis focuses on a method first developed by von Helmholtz at the end of the nineteenth century, called prism adaptation (PA) (von Helmholtz 1867). In this paradigm, participants wear prism glasses that bend light, and so optically displace the visual field, for example by 10° to the right. When participants perform visually guided motor tasks (e.g. pointing at targets) while wearing the prisms, at first, they make systematic rightward errors (owing to the optical displacement), but participants learn rapidly from the error feedback to correct their movements on subsequent trials and regain normal accuracy (i.e. they adapt). When the prisms are removed post-adaptation, individuals then

make errors in the opposite direction, i.e. a leftward after-effect (AE), which reflects the temporary persistence of some of the compensatory mechanisms engaged during the adaptation.

Several features of how prism AEs generalise or transfer beyond the specifically trained context make it an interesting paradigm to investigate. In healthy controls, prism AEs tend to generalise at least partially across space (Bedford 1989; Bedford 1993; Redding & Wallace 2006a). This contrasts with visuomotor rotation, for instance, where effects drop off sharply with distance from the trained target location (Krakauer *et al.* 2000). Pointing during prism exposure is typically aimed at lateral targets under speeded conditions, whereas prism AEs are often measured at a central (untrained) location, with accuracy emphasised over speed. With this procedure, there is, therefore, a change in task context between prism exposure and prism AE measurement conditions, such that the AE measure intrinsically captures elements of generalisation/transfer, at least with respect to task changes in reach trajectory, reaching movement speed and target location. Prism AEs are measurable in at least three different modalities, visual, proprioceptive, and motor, which appear to follow different dynamics (Harris 1963; Redding & Wallace 2001; Hatada *et al.* 2006a; Hatada *et al.* 2006b; Hatada *et al.* 2006c). It has also been claimed that prism AEs can transfer to untrained visuospatial tasks (e.g. line bisection task, greyscales task), although these effects in young healthy volunteers appear to occur only with left-shifting (not right-shifting) prisms and tend to be quite small and variable (Colent *et al.* 2000; Michel *et al.* 2003; Loftus *et al.* 2009; Goedert *et al.* 2010; Schintu *et al.* 2014; Martín-Arévalo *et al.* 2016; Striemer *et al.* 2016; Schintu *et al.* 2017).

A stronger evidence base in patients has shown that the after-effects of PA can transfer to improve cognitive deficits in spatial neglect after right hemisphere brain

damage (Rossetti *et al.* 1998; Frassinetti *et al.* 2002; Serino *et al.* 2009; Barrett *et al.* 2012; Rode *et al.* 2015; O’Shea *et al.* 2017). Spatial neglect is a complex multi-faced neurological syndrome that affects approximately half of right-hemisphere stroke patients (Stone *et al.* 1993; Karnath *et al.* 2004; Ringman *et al.* 2004; Milner & McIntosh 2005; Verdon *et al.* 2009). Its core defining feature is the inability for patients to voluntarily orient attention to the contralesional left side of space (Heilman & Van Den Abell 1980; Heilman *et al.* 2002). AEs have been shown to transfer to a broad range of untrained sensory and cognitive domains in neglect, including, for example, postural control, occulo-motor exploration, dichotic listening and mental imagery (for review, see Jacquin-Courtois *et al.* 2013). This peculiar generalisation pattern makes PA a very attractive tool for stroke rehabilitation where one of the greatest challenges is to design interventions that can benefit patients beyond the restricted context of the trained task (Langhorne *et al.* 2011). However, the current limitations of PA-based neglect rehabilitation include the relatively short-lasting duration of the reported therapeutic effects (Rossetti *et al.* 1998; Farnè *et al.* 2002; Pisella *et al.* 2002; Jacquin-Courtois *et al.* 2008), and the large inter-patient variability in the response to PA (Barrett *et al.* 2012; Bowen *et al.* 2013). This is partly due to a poor fundamental understanding of the neural mechanisms involved. Thus, the main objective of this thesis is to better characterise the neuro-computational mechanisms of sensorimotor memory formation during PA in the healthy brain, in order to maximise AE persistence and improve PA-based neglect rehabilitation.

The first step of this bench-to-bedside translational work is to define a theoretical framework within which to investigate PA. What features should a satisfying theoretical account of PA behaviour have? An ideal account would provide: 1) mechanistic explanations, that are 2) biologically plausible, and 3) can generate quantitative behavioural predictions, 4) about the effect of a range of factors, such as experimental

task manipulations (e.g. modality, quality and timing of sensory feedback, gradual versus abrupt perturbation onset, etc.), psychological variables (e.g. internal state estimates of limb position, sensory uncertainty, prior knowledge of the perturbation, etc.), and neural state effects (e.g. change in neural excitability in specific brain region owing to lesion, drug or brain stimulation interventions). Although some of these points are outside the scope of this thesis, we undertake to use a theoretical framework that meets as many of these criteria as possible.

In the present chapter, we discuss the potential of PA interventions for neglect rehabilitation, and argue that a better fundamental understanding of PA is needed to improve efficacy of such interventions. Next, we outline the current (descriptive psychological) model of PA that is predominant in the literature on healthy individuals, patients and animal studies. We also highlight the brain regions implicated in PA by studies conceived within this framework. Finally, we make the case that a computational characterisation of PA behaviour offers advantages over this traditional functional descriptive approach, and argue the need for an integrated neuro-computational account to further advance understanding within this field.

1.2 Prism adaptation procedures

Since the primary focus of this thesis is on PA experiments, we first give a practical guide on how such studies are typically performed. This will provide the reader with the necessary background to engage with the PA literature and understand the theoretical discussion developed in the following sections.

PA experiments usually include at least three phases, as follows: 1) *pre-adaptation*

phase: baseline measure of individuals' visuomotor task performance; 2) *prism exposure phase*: visuomotor tasks are performed during exposure to the prismatic shift. This optical shift causes performance errors, and individuals learn to correct their errors gradually and to compensate for the optical shift (i.e. they adapt); 3) *post-adaptation phase*: baseline tests are repeated, and changes in performance (post - pre) provide a measure of AEs (Kornheiser 1976) (Figure 1.1).

In the studies we will consider, participants are typically asked to make pointing movements during each of the three phases. Visual feedback of individuals' trial-by-trial reach endpoints can either be provided (closed-loop pointing, CLP) or deprived (open-loop pointing, OLP). During prism exposure (phase 2), endpoint errors are deviated initially in the direction of the optical displacement on both of these types of pointing. Performance change relative to baseline (during - pre) during this initial phase is often described as the direct effect of prisms. If visual feedback of endpoint errors is provided during prism exposure (i.e. closed-loop pointing), individuals tend to reduce these errors progressively to regain baseline accuracy (i.e. they adapt) (Figure 1.1B-C). Prism AEs are measured by asking participants to point again after removal of the prisms. If sufficient practice has occurred during prism exposure, performance will now be deviated in the direction opposite to the prismatic shift. If visual feedback of these (now leftward) reach endpoints is provided (i.e. closed-loop pointing), performance errors will once again be corrected rapidly (now in the opposite direction) to re-gain baseline accuracy (Figure 1.1C). This de-adaptation (or washout) can be limited by depriving visual feedback (i.e. using open-loop pointing) (Figure 1.1D). Prism AEs measured using open-loop pointing can persist even after closed-loop performance has returned to baseline levels (Inoue *et al.* 2015).




Figure 1.1: Prism adaptation. **A.** By bending light, prism lenses displace the visual field in a direction determined by the prism structure. Here for example, light is displaced laterally, by 10° to the right. Hence, a central dot when viewed through this prism is (mis)perceived to be located 10° to the right of its true position. **B.** Schematic representation of the typical experimental phases of PA. At baseline (1), pointing is accurate. During prism exposure, participants make rightward pointing errors initially (early phase, 1), but learn gradually from trial-by-trial error feedback to correct these errors and re-gain close-to-baseline pointing accuracy (late phase, 2). Consequent leftward prism AEs (errors) are measurable post-adaptation once the glasses have been removed (3). **C.** Canonical pattern of performance errors during PA. Plot shows reach endpoint error (y-axis) as a function of trial number (x-axis) during closed-loop pointing (i.e. with visual feedback). This typical pattern consists of baseline accuracy (i.e. mean error centred on zero) (1), followed by rightward errors (in the direction of the prismatic shift) that decrease gradually across prism exposure trials (2), followed by leftward errors (in the direction opposite the prismatic shift) after removal of the prism goggles (3). **D.** Three tests commonly used in the PA literature to quantify prism AEs. The sum of visual and proprioceptive AEs immediately after prism exposure has been shown to equal the magnitude of AE quantified by open-loop pointing, which is therefore known as the total visumotor shift (Templeton *et al.* 1974).

Two other tests are commonly contrasted across the baseline and post-exposure test phases, to measure modality-specific prism AEs: visual straight-ahead judgements and proprioceptive straight-ahead pointing (Redding & Wallace 1996; Redding & Wallace 2001; Hatada *et al.* 2006b; Hatada *et al.* 2006c). The idea is to measure the different components hypothesised to underlie PA: visual, proprioceptive and motor changes (Harris 1963). The visual straight-ahead judgement requires participants to report verbally when a visual stimulus moving laterally across their visual field is per-

ceived as being straight-ahead of their nose (Figure 1.1D). This measure is thought to rely on eye-head coordination, and any post-exposure change in accuracy is usually interpreted as a visual AE (Hatada *et al.* 2006a). Proprioceptive straight-ahead pointing requires blindfolded participants to point to the position in space they perceive to be straight-ahead of their nose. It is thought to provide a measure of head-hand coordination, and post-exposure shifts in this measure are interpreted as a proprioceptive AE (Harris 1963; Templeton *et al.* 1974). Several authors have reported that immediately after PA, the sum of the visual and proprioceptive AEs – measured by visual straight-ahead judgement and proprioceptive straight-ahead pointing, respectively – equals the magnitude of AEs assayed using open-loop pointing (Hay & Pick Jr 1966; Templeton *et al.* 1974; Redding & Wallace 1996; Hatada *et al.* 2006c). This finding is relatively intuitive given that both vision of target location at onset and proprioceptive and motor feedback during the reach trajectory contribute to motor performance on open-loop pointing. Open-loop pointing quantifies the combined contribution of both factors influencing hand-eye coordination (i.e. the total visuomotor shift), while visual straight-ahead judgement and proprioceptive straight-ahead pointing each measure individual components (eye-head and head-hand respectively). Studies vary in whether they measure just the total visuomotor shift, or the visual or proprioceptive AEs, or investigate the relationship between all three.

1.3 Rehabilitation potential for spatial neglect

Spatial neglect is a complex neurological syndrome occurring mostly in right hemisphere stroke patients (Stone *et al.* 1993; Karnath *et al.* 2004; Milner & McIntosh 2005; Verdon *et al.* 2009), and characterised by the failure to report, respond or orient to novel or meaningful stimuli presented to the contralesional (left) side, when this

failure cannot be attributed to either sensory or motor deficits (Heilman & Van Den Abell 1980; Heilman *et al.* 2002). It is a strongly debilitating condition that limits the overall outcome of rehabilitation by adversely affecting mobility, discharge destination, length of hospital stay and independence in self-care skills (Barer 1990; Jehkonen *et al.* 2006; Patel *et al.* 2007; Di Monaco *et al.* 2011). Therapeutic approaches for spatial neglect can be categorised in three groups: top-down, bottom-up, and neuromodulatory strategies (Azouvi *et al.* 2016; Rode *et al.* 2017). Top-down techniques are by far the most commonly used therapeutic approach (Bowen *et al.* 2007; Bowen *et al.* 2013) and refer to empirical methods used to train patients to generate voluntary shifts of attention to the neglected side of space. For example, visual scanning training (Weinberg *et al.* 1977; Weinberg *et al.* 1979), which forms the basis of neglect therapy delivered in most rehabilitation units (Barrett *et al.* 2006), consists of prompting patients to look and orient leftwards, using auditory, visual, or other cues. It is extremely time-consuming¹ and may be inappropriate from a theoretical perspective, due to the lack of insight on their condition (*anosognosia*) that neglect patients often exhibit (Appelros *et al.* 2002). The other two approaches (bottom-up and neuromodulatory techniques) try to remedy this latter limitation by by-passing the need for a voluntary effort of patients to overcome their pathological behavioural bias. Bottom-up approaches are based on manipulations of the sensory environment to induce involuntary leftward behavioural shifts. Neuromodulatory techniques use brain stimulation tools (e.g. transcranial magnetic stimulation, transcranial direct current stimulation; [chapter 2](#)) to influence visuospatial processing at the physiological level, so as to re-balance the representation of left and right space (for an example of such modulation in the healthy brain, see [Appendix B](#)). The paradigm of interest, prism adaptation, is therefore a bottom-up technique.

¹In the United States, patients suffering from spatial neglect typically spend 1 hour daily on this training alone, for at least 4 consecutive weeks (Barrett *et al.* 2012).

Standard PA interventions (Rossetti *et al.* 1998) require patients to wear right-shifting prisms while making visually guided pointing movements (section 1.2; [Figure 1.1](#)). It is the induction of a leftward AE during the adaptation that is thought to induce clinical improvements (Serino *et al.* 2006; Rode *et al.* 2015). The advantages of this technique over the traditional top-down approach mainly include its 1) simplicity and 2) generalisation pattern. Stroke survivors undergoing PA interventions wear prisms for only a short period of time (typically 15-20 minutes per session) and are free to take part in other therapies or medical procedures during the rest of the day. Thus, compared to the hours of training involved in most top-down rehabilitation approaches, PA training is less time-consuming. Since Rossetti *et al.* (1998)'s seminal experiment, most studies reported improvement on visuo-motor tasks such as line bisection, cancellation or drawing following PA in neglect patients (for review, see Striemer & Danckert 2010; Barrett *et al.* 2012; Newport & Schenk 2012; Bowen *et al.* 2013; Jacquin-Courtois *et al.* 2013). Additionally, some studies reported beneficial effects of PA on untrained cognitive and sensory domains such as reading (Serino *et al.* 2009), mental imagery (Rode *et al.* 2001; Rossetti *et al.* 2004), wheelchair driving (Jacquin-Courtois *et al.* 2008; Watanabe & Amimoto 2010), dichotic listening (Farnè *et al.* 2002; Jacquin-Courtois *et al.* 2010) and measures of functional independence (Mizuno *et al.* 2011). These reports illustrate the peculiar 'expansion' of prism AEs beyond the training conditions, one of the key objectives of post-stroke functional rehabilitation.

Despite these apparent advantages, we find that 20 years after the introduction of PA as a neglect rehabilitation tool (Rossetti *et al.* 1998), only a small number of inpatient centres use this approach, mainly within research protocols. Why is this technique not used more broadly? We believe that this is due mainly to a lack of fundamental understanding of the mechanisms involved (in both spatial neglect and

prism adaptation) (Striemer & Danckert 2010; Barrett *et al.* 2012), that is responsible for unexplained variance in the clinical response to PA interventions. The evidence base for the clinical efficacy of PA relies mainly on studies including single cases or small groups of patients (e.g. Rossetti *et al.* 1998; Farnè *et al.* 2002; Frassinetti *et al.* 2002; Serino *et al.* 2006; Shiraishi *et al.* 2008; Saevarsson *et al.* 2009; Vangkilde & Habekost 2010; Lâdavass *et al.* 2011) and shows large inter-patient variability in the clinical response to PA². It is not to say that case studies have no scientific value, but rather that, just as the neglect syndrome is heterogeneous and highly variable in presentation (Azouvi *et al.* 2002), the influence of PA on neglect could also be heterogeneous and variable across patients (Striemer & Danckert 2010). This might explain why, with over 90 patients tested in total, systematic randomised clinical trials (RCTs) demonstrated only small to no definite clinical benefit at the group level compared to control interventions (Nys *et al.* 2008; Serino *et al.* 2009; Turton *et al.* 2010; Mizuno *et al.* 2011; Rode *et al.* 2015; Ten Brink *et al.* 2017).

To date, the mechanisms responsible for neglect improvements following PA remain largely unknown (Striemer & Danckert 2010; Barrett *et al.* 2012; Newport & Schenk 2012). However, it appears that across patients, greater sensorimotor AE is associated with greater therapeutic response to PA (Serino *et al.* 2006; Sarri *et al.* 2008; Rode *et al.* 2015). Thus, translation of basic science knowledge of the neural and psychological mechanisms involved in the formation of an AE memory trace during PA might improve rehabilitation efficacy, and help in identifying patients that are likely to respond to PA. The following sections review the current dominant view on the mechanisms involved in PA in healthy individuals.

²Approximately 50-75% of tested patients respond to PA.

1.4 The traditional dual-process framework: strategic control versus spatial realignment

1.4.1 Theoretical framework

For the past forty years, the large majority of studies investigating the neural mechanisms underlying PA have interpreted their results within a theoretical framework that distinguishes two learning processes that contribute differentially to the direct effects (i.e. error correction) and AEs of prisms. This framework posits that PA recruits two distinct functional mechanisms: a rapid process of error reduction that reflects strategic adjustments in motor control, and a slower process, so-called 'true' sensorimotor adaptation, thought to reflect the spatial realignment of motor, proprioceptive and visual coordinate reference frames (for review, see Redding & Wallace 2002; Redding *et al.* 2005; Redding & Wallace 2006b). The name given to these two processes has varied slightly across studies and over time, but the core idea of a distinction between a fast strategic component and a slower 'true' sensorimotor realignment has remained consistent. Here we will adopt Redding and Wallace's (Redding *et al.* 2005; Redding & Wallace 2006b) latest terminology to describe this dual-process theoretical framework, and review the evidence for functional and anatomical dissociations associated with these two processes.

Within this framework, the constituent processes driving PA are described as follows. *Strategic control* is a set of processes that guide everyday adaptive motor behaviours. For example, to reach for a cup, depending on the sensory information available, it is necessary to first select the appropriate reference frame to code the target location (e.g. visual-motor, proprioceptive-motor) and guide the appropriate reach-to-grasp command. The process of setting the reference frame in relation to the desired action has been called 'calibration' (Redding & Wallace 2002; Redding




Figure 1.2: Traditional dual-process framework. The traditional theoretical framework posits that error correction during PA relies on two processes. Strategic control refers to the calibration of individuals' task workspace around the task relevant objects. It is described as a rapid process that intervenes early during prism exposure but contributes to a lesser extent to prism AEs. Conversely, spatial realignment is described as developing more slowly during prism exposure and is thought to be responsible for the prism AEs. The term describes a process of bringing the different sensorimotor coordinate frames (visual-motor, proprioceptive-motor) into alignment with each other.

et al. 2005). Strategic control also requires selecting the region of extrapersonal space most relevant to the on-going task (e.g. the shelf where the cup is located), and is therefore described as being closely related to spatial attention. If a reach-to-grasp action is unsuccessful, the action may have to be 'recalibrated' by various means. For example, (while wearing prisms) one may direct his/her reaching movement more towards the side of the cup, so as to reduce the previous motor error. This mix of partly automatic, partly conscious processes is thought to contribute predominantly to the rapid error correction that occurs during the initial phase of prism exposure, but to contribute to a lesser extent to prism AEs (Weiner *et al.* 1983; Redding *et al.* 2005; Redding & Wallace 2006b; Aimola *et al.* 2012) (Figure 1.2). It is important to note that, within this framework, strategic motor control adjustments do not require the linkage between different coordinate frames to be reconfigured. Instead, motor performance can be improved simply by changing the motor command issued to reach

the target. An analogy commonly used to illustrate this idea is to consider a rifleman with misaligned telescopic sight. Suppose that the scope of the rifle is misaligned with the barrel, so that the marksman misses the target systematically by 10° to the right. Simply re-aiming 10° to the left of the target would allow the marksman to hit his target without having to realign the scope with the barrel. Strategic control refers to this process of quickly recalibrating the reference frame to meet the task objectives. Successful movements driven by such a process should therefore be specific to the trained task context and should not generalise to other contexts (e.g. other movements, other spatial locations, etc.).

By contrast, *spatial realignment* refers to the process of adjusting for constant differences in spatial coordinates between multiple sensorimotor coordinate systems. In reference to the previous analogy, this would be equivalent to realigning the barrel with the scope of the rifle. The traditional theoretical framework often posits that accurate pointing movements require the correct alignment of two reference frames: an eye-head visual-motor system (measured with the visual straight-ahead test) and a head-hand proprioceptive-motor system (measured with straight-ahead pointing) (Redding & Wallace 2002; Redding *et al.* 2005). Because the prismatic shift displaces the visual-motor reference frame only, compensatory shifts in the visual-motor and/or proprioceptive-motor systems are required in order to re-align the two reference frames and regain correct eye-hand coordination during prism exposure (Templeton *et al.* 1974). The temporary carry-over of these adaptive shifts after removal of prisms is thought to be predominantly responsible for the after-effects of PA.

1.4.2 Behavioural predictions

In summary, the core theses of the traditional dual-process framework are that: 1) strategic recalibration is a cognitively demanding process, that drives error correc-

tion early during prism exposure, but contributes little to prism AEs; 2) spatial realignment is an automatic process, that develops more gradually, and is mainly responsible for prism AEs; 3) these two processes operate relatively independently from one another (Figure 1.2). Within this theoretical framework, closed-loop performance during prism exposure reflects both the strategic control and spatial realignment components, whereas prism AEs (measured via open-loop pointing, visual straight-ahead judgement or proprioceptive straight-ahead pointing) provide a measure of spatial realignment only. Impaired strategic control would therefore affect error correction during prism exposure, but would not directly affect prism AEs. By contrast, impaired spatial realignment would affect both direct- and after- effects of prisms. In the following two sections (1.4.3 and 1.4.4), we review some of the main empirical evidence that supports this theoretical model (for in depth review, see Redding & Wallace 2002; Redding *et al.* 2005; Redding & Wallace 2006b) and summarise attempts to localise the neural correlates of these two proposed processes.

1.4.3 Strategic control

Within the traditional framework, strategic control refers to the process of quickly recalibrating the reference frame to (re)code the location of the target. It is thought to contribute mainly to the early phase of prism exposure. Consistent with this view, several studies have reported behavioural markers of a learning component engaged during the early phase of prism exposure, but saturating quickly, and contributing to a lesser extent to the AE. For example, in healthy individuals, during the initial phase of prism exposure, when errors are large, the time gap between target foveation and onset of the pointing movement increases transiently, but returns to normal within about 10 trials, as the endpoint error magnitude is rapidly reduced (Rossetti *et al.* 1993). A similarly rapid time course was observed in trial-by-trial corrections of the initial acceleration phase of the reach trajectory during the first 10

prism exposure trials (O'Shea *et al.* 2014). It has also been shown that imposing a cognitive load (mental arithmetic) during prism exposure disrupts participants' ability to correct pointing errors while wearing prisms (Redding *et al.* 1992). None of these behavioural phenomena have been shown to relate quantitatively to measures of prism AEs. Taken together, these results implicate a strategic learning component that contributes to error correction during early prism exposure.

The neural substrates associated with this 'strategic control' component of PA have been inferred from brain lesions that occur concurrent to impaired error reduction during prism exposure, but spared AEs (Welch & Goldstein 1972; Weiner *et al.* 1983; Canavan *et al.* 1990; Pisella *et al.* 2004; Newport & Jackson 2006; Fernandez-Ruiz *et al.* 2007) (Table 1.1). Such studies have converged on a crucial involvement of the cerebral cortex in strategic control (Welch & Goldstein 1972; Weintraub & Mesulam 1987; Canavan *et al.* 1990), more specifically, the posterior parietal cortex (Pisella *et al.* 2004; Newport & Jackson 2006). The functional specificity implied by this behavioural pattern is debatable given the variety of neurological impairments that can lead to this pattern of performance. For example, temporal lesions (Canavan *et al.* 1990), spinocerebellar ataxia type 2 (Fernandez-Ruiz *et al.* 2007), and spatial neglect (Aimola *et al.* 2012) have also been associated with impaired error reduction but intact prism AEs.

Study	Shift	Hand	Etiology	Impaired SC	Impaired SR
<i>Cerebellar lesions</i>					
Weiner <i>et al.</i> 1983	R	R/L	Degenerative (n=5), infarct (n=1), drug induced (n=1)		✓
Baizer <i>et al.</i> 1999	L	L/R	Aspiration lesion: left posterior cerebellum, including caudal parts (n=1)		✓

	L	L/R	As above but not including the caudal parts (n=4)	o	o
Pisella <i>et al.</i> 2005	L	R	Left anterior cerebellar stroke (n=1)		✓
	L	L			✓
	R	R		o	o
	R	L		o	o
Fernandez-Ruiz <i>et al.</i> 2007	R	R	Spinocerebellar ataxia type 2 (SCA2) (n=43)	✓	
Calzolari <i>et al.</i> 2015	R	R	Left lateral cerebellar stroke + bilateral occipital lesion (n=1)		✓
Hanajima <i>et al.</i> 2015	R/L	R	Cerebellar cortical atrophy (n=13)		✓
Hashimoto <i>et al.</i> 2015	R	R	Degenerative cerebellar diseases (n=77)		✓
<i>Parietal lesions</i>					
Pisella <i>et al.</i> 2004	L	R	Bilateral PPC stroke (n=1)	✓	
Newport & Jackson 2006	R	R	Bilateral PPC atrophy,	✓	
	R	L	predominantly in the	✓	(✓)
Newport <i>et al.</i> 2006	R	R	left hemisphere (n=1)	✓	(✓)
	R	L		✓	(✓)
<i>Other</i>					
Welch & Goldstein 1972	L	R	Various brain lesions (n=18)	✓	
Weiner <i>et al.</i> 1983	R	R/L	Left hemisphere lesion (n=15)	o	o
	R	R/L	Right hemisphere lesion (n=14)	✓	
	R	R/L	Parkinson's disease (n=10)	(✓)	
Canavan <i>et al.</i> 1990	?	R	Frontal lesions (n=9)	✓	
	?	R	Left temporal lobectomy (n=9)	o	o
	?	R	Right temporal lobectomy (n=10)	✓	
	?	R	Early Parkinson's disease (n=19)	✓	
Stern <i>et al.</i> 1988	L	R	Parkinson's disease (n=14)		✓
Fernandez-Ruiz <i>et al.</i> 2003	L	R	Parkinson's disease (n=35)		✓
Aimola <i>et al.</i> 2012	R	R	Spatial neglect (n=11)	✓	

Table 1.1: Overview of the lesion studies of prism adaptation. This table includes the main lesion studies of prism adaptation in primates and humans. It is not meant to be exhaustive. Check mark indicate impaired strategic control or spatial realignment. The 'o' indicates preserved strategic control and spatial realignment. R = right; L = left; PPC = Posterior Parietal Cortex; SC = strategic control; SR = spatial realignment

In the healthy brain, several groups have used neuroimaging techniques to try and identify brain regions involved in strategic control during PA. Typically, these studies have contrasted the amplitude of neural activity in the 'early' versus 'late' phases of prism exposure and interpreted areas activated by this contrast (i.e. early > late) as neural correlates of the strategic control component of PA (Table 1.2). Consistent with lesion studies, activity in the posterior parietal cortex has been reliably reported by whole brain neuroimaging studies, but other brain regions such as primary motor cortex, anterior cingulate cortex or cerebellum have also been found to be preferentially activated during the early versus late phase of prism exposure (Clower & Boussaoud 2000; Danckert *et al.* 2008; Luauté *et al.* 2009; Küper *et al.* 2014) (Table 1.2). Aside from functional localisation in this way, these neuroimaging studies have said little about how these identified brain regions are thought to implement the process of strategic control.

Study	Shift	Hand	Area	Side	MNI coord.		
					<i>x</i>	<i>y</i>	<i>z</i>
Clower & Boussaoud 2000	L/R	R	IPS	L	-51	-49	44
			M1	L	-36	-37	-16
Danckert <i>et al.</i> 2008	R	R	ACC	R	2	-1	49
			ant. IPS	L	-39	-40	51
			med. Cer	R	1	-62	-4
Luaute <i>et al.</i> 2009	L	R	ant. Cer	R	24	-38	-32
			IPS/SPL	R	40	-54	66
			IPS/IPL	L	-30	-68	40
Küper <i>et al.</i> 2014*	R	R	post. Cer (VIII)	R	31	-57	-52
			post. Cer (IX)	R	6	-62	-42
			cDN	R	13	-67	-32
					18	-66	-36

Table 1.2: Neuroimaging correlates of strategic control. * Study including the cerebellum and dentate nucleus only. R = right; L = left; ACC = Anterior Cingulate Cortex; cDN = caudate Dentate Nucleus; Cer = Cerebellum; IPL = Inferior Parietal Lobule; IPS = Inferior Parietal Sulcus; M1 = Primary Motor Cortex; SPL = Superior Parietal Lobule.

1.4.4 Spatial realignment

Within the traditional dual process framework, spatial realignment refers to the process of shifting the visual-motor and/or proprioceptive-motor systems in order to realign these two coordinate frames and regain accurate hand-eye coordination. This process is believed to develop slowly during prism exposure, and to be responsible almost entirely for the prism AEs (Redding & Wallace 2002; Redding *et al.* 2005; Redding & Wallace 2006b). One study in healthy individuals reported a putative kinematic signature of sensorimotor realignment: gradual correction of the terminal (deceleration) phase of the pointing trajectory during prism exposure. This corrective process unfolded slowly during prism exposure and the magnitude of correction of this kinematic error correlated with the magnitude of prism AEs (O'Shea *et al.* 2014). Endpoint error appears not to be necessary for spatial realignment to occur, as AEs can be observed in the absence of measurable reach endpoint errors if the

prismatic shift is introduced gradually (Dewar 1971; Michel *et al.* 2007; Hanajima *et al.* 2015). Instead, it has been suggested that the conflict between the expected (feedforward predicted) and observed (feedback measured) position of the hand is the learning signal for spatial realignment (Redding & Wallace 2006b). Support for this claim comes mainly from the finding that the magnitude of visual and proprioceptive AEs (measured by the visual straight-ahead judgement and proprioceptive straight-ahead pointing, respectively) depends upon the sensory information available in-flight during the reach trajectory when individuals are wearing prisms (Redding & Wallace 1996; Redding & Wallace 2001). If both proprioceptive and visual feedback is available, proprioceptive AEs tend to be greater than visual AEs. The opposite is true (i.e. greater visual than proprioceptive AEs) if only proprioceptive feedback is available in-flight (Redding & Wallace 1996; Redding & Wallace 2001). This suggests that the modality in which the discrepancy between the predicted and observed hand location is sensed determines which coordinate reference frame is preferentially re-aligned.

Study	Shift	Hand	Area	Side	MNI coord.		
					<i>x</i>	<i>y</i>	<i>z</i>
Luauté <i>et al.</i> 2009	L	R	STS/STG	R	58	-12	-12
					70	-20	4
			STS/STG	L	-60	-26	4
Chapman <i>et al.</i> 2010	L	(R)	post. Cer	R	16	-58	-50
			ANG	R	36	-78	30
			ant. IPL	R	42	-50	46
Küper <i>et al.</i> 2014*	R	R	post. Cer (VI)	R	3	-60	-19

Table 1.3: Neuroimaging correlates of spatial realignment * Study including the cerebellum and dentate nucleus only. R = right; L = left; ANG = Angular Gyrus; Cer = Cerebellum; IPL = Inferior Parietal Lobule; M1 = Primary Motor Cortex; STS/STG = Superior Temporal Sulcus/ Superior Temporal Gyrus.

Several studies have reported evidence of impaired spatial realignment following

cerebellar lesions in humans (Weiner *et al.* 1983; Martin *et al.* 1996; Pisella *et al.* 2005; Calzolari *et al.* 2015; Hanajima *et al.* 2015) (Table 1.1) and non-human primates (Baizer *et al.* 1999). The typical behavioural pattern is reduced error correction during prism exposure, combined with decreased or absent prism AEs. The contribution of the precise cerebellar sub-regions is still unclear, as evidence of impaired spatial realignment has been found after anterior (Pisella *et al.* 2005; Calzolari *et al.* 2015) and posterior (Martin *et al.* 1996; Baizer *et al.* 1999) lesions to cerebellum. To our knowledge, only three neuroimaging studies have investigated the pattern of functional brain activity associated with spatial realignment (Luauté *et al.* 2009; Chapman *et al.* 2010; Küper *et al.* 2014). They did so by investigating brain regions that were more active during the later stage of prism exposure compared to the early stage (late > early). Posterior cerebellar activity was reported in two of these studies (Chapman *et al.* 2010; Küper *et al.* 2014), but other regions such as the superior temporal gyrus and angular gyrus were also activated (Luauté *et al.* 2009; Chapman *et al.* 2010) (Table 1.3).

1.4.5 Why is this framework not satisfying?

The traditional dual-process theoretical framework offers an account of various behavioural dissociations observed in healthy individuals and neurological patients (for review, see Redding & Wallace 2002; Redding *et al.* 2005; Redding & Wallace 2006b). The main functional insight provided by this framework has been to distinguish two psychological processes (strategic control, spatial realignment) that combine to explain behaviour, but seem to operate with a certain degree of independence from one another. However, studies conceived within this theoretical framework do not offer a mechanistic explanation of how PA behaviour arises, and attempts to localise the neural circuits underlying these two processes have yielded heterogeneous results (Table 1.1, 1.2, 1.3). In addition, the information processing operations executed by

the identified neural components remain largely unknown.

The traditional dual-process theory suffers from a pervasive problem in cognitive neuroscience – how to bridge the conceptual gap between cognitive psychological level descriptions of behaviour and biologically plausible descriptions of neural circuit functioning? Computational models offer a potential bridge, as they provide a common currency (algorithms, mathematical functions) in which to describe both information processing and neural circuit mechanisms that could implement such functions. We contend that, in order to advance cognitive neuroscience explanation of the causal brain-behaviour dynamics that underwrite PA, the traditional dual-process framework needs to be replaced with a re-conceptualisation at the algorithmic level. By 'algorithmic' we mean a level of description that sets out clearly the (mathematical) rules and operations (functions) required to execute PA behaviour. It is obvious that, at the cellular and neural circuit level, cognitive concepts like 'strategic control' or 'spatial realignment' have no explanatory value, since brain circuits are computing information, and the explanatory task is to provide an account of how these computations implement psychological functions. Explanatory progress therefore requires a conceptual advance: a theoretical framework that decomposes PA behaviour into the underlying algorithms required to implement it (for a related argument, see Krakauer *et al.* 2017). Re-conceptualising PA in terms of its constituent algorithms offers objective mathematical description, as opposed to qualitative description present in the traditional dual-process psychological framework. This greater precision helps avoid confusion related to terminology. It also allows for quantitative experimental predictions. We will return to these points in the next section. More generally, recasting any behaviour or cognitive process in algorithmic terms helps advance the field towards a re-defined taxonomy of cognitive processes, one grounded in the recognition that the same computations (and brain circuits) might contribute to diverse behavioural

phenomena depending on the constraints of the task. This offers a way to move beyond accounts of brain activations in terms of 'area X is involved in cognitive function Y (e.g. attention)' to 'computation Y is implemented within circuit X and engaged during tasks a, b, c'. This distinct conceptual approach could, for example, help to explain why prism AEs transfer from pointing tasks to cognitive domains in neglect patients (Rossetti *et al.* 1998; Sumitani *et al.* 2007).

In the following section, we propose an algorithmic decomposition of the computations required for PA, by leveraging insights from the field of computational neuroscience applied to sensorimotor control. We argue that this decomposition, which is not incompatible with the traditional psychological approach, offers a more useful theoretical framework, in terms of quantitative precision, explanation, and plausible neural implementation. Because it is directly relevant to most of the work presented in this thesis (see [chapter 4](#)), only the state-space modelling framework is evoked in this introductory chapter. See [Appendix A](#) for a review of an other possible computational framework, Bayesian Decision Theory (Körding & Wolpert 2006).

1.5 An alternative theoretical framework: multiple timescale models

State-space modelling has provided insights into the hidden internal state learning and forgetting dynamics that contribute to overt behaviour during sensorimotor adaptation (Smith *et al.* 2006a). Trial-by-trial error correction has been explained as the output of multiple internal adaptive systems that learn and forget on different timescales. The key idea is that these systems compete to learn from performance error and that the sum of their states provides an inverse estimate of the perturbation that is used to correct motor performance ([Figure 1.3A](#)). A commonly used simplified

model posits that the multitude of possible timescales can be approximated by a fast system, which learns and forgets rapidly, and a slow system, which learns and forgets more slowly (Smith *et al.* 2006a). The model posits that early during adaptation, error correction is mainly dominated by the fast system, which saturates quickly before decaying back to baseline. By contrast, the slow system develops more gradually over extended practice and accounts for most of the net adaptation at later stages of exposure to the perturbation. This 'two-state' framework is able to reproduce and explain a range of behavioural phenomena in adaptation, including apparently counterintuitive effects such as spontaneous recovery - the re-appearance of AEs after a brief period of washout (Smith *et al.* 2006a; Zarahn *et al.* 2008; Lee & Schweighofer 2009). The model offers an efficient way to extract (hidden, internal state-driven) temporal dynamics of learning processes that drive adaptation and to generate and test precise quantitative predictions about the relative contributions of these fast versus slow systems to behaviour. For example, it has been shown that the level of adaptation reached by the slow system during force-field adaptation, rather than the overall performance improvement, predicts the amount of long-term retention (Joiner & Smith 2008).

Because they are not restricted to two learning systems (Körding *et al.* 2007; Inoue *et al.* 2015; Kim *et al.* 2015), multiple timescale models can also be used to ask what number of processes best explains adaptation behaviour. For example, a recent PA study asked how many systems were needed to explain the immediate and long-term retention dynamics of prism AEs (Inoue *et al.* 2015). During prism exposure, error-dependent learning (CLP trial accuracy) asymptotes quickly after the initial error correction phase, i.e. once pointing accuracy has been restored to near-baseline levels, it remains at the same level for the rest of the exposure period. However, even once CLP behaviour has plateaued, the subsequent AE dynamics (on OLP trials)

continue to change, depending on the duration of the preceding prism exposure. Typically, the longer the prism exposure, the more stable the subsequent AEs (i.e. less decay), reflecting consolidation of the acquired adaptation. To account for this phenomenon, Inoue and colleagues proposed a third 'ultraslow' system, required to explain behaviour only when prism exposure was prolonged (500 trials) (Figure 1.3B).




Figure 1.3: The three timescales state-space model for prism adaptation. **A.** Each of the fast, slow and ultraslow systems are described by a pair of free parameters: a retention factor A , that describes the amount of decay occurring between trials ($0 < A_f < A_s < A_{us} < 1$), and a learning rate B , that describes the fraction of performance error being incorporated into that system's state on each trial ($0 < B_{us} < B_s < B_f < 1$). The sum of the states of the three systems produces an inverse estimate of the perturbation on a trial-by-trial basis, which is used to correct motor output. On any trial, performance error therefore corresponds to the sum of the prismatic shift and the state of the three systems. We have chosen here to illustrate state-space models with a three timescales model because of its relevance for PA (Inoue *et al.* 2015), but in principle any number of systems could be posited. **B.** Simulation of PA data with a three timescales model. When the prismatic shift (in blue) is introduced, the three systems (dotted yellow, green and purple lines) learn at three different rates (set by their respective learning rates B). The sum of the states of the three systems is added to the magnitude of the prismatic shift to reduce performance error (in red) during prism exposure.

(Legend continued on next page)

If the AEs are probed after 30 exposure trials (learning rate is set to 0 for all three systems during AE measurements, i.e. OLP), the memory trace decays rapidly, because it is largely dominated by the fast system, which has the lowest retention factor. However, if prism exposure continues, the contribution of the fast system gradually decreases. As a result, AEs are more stable after 150 trials. Finally, after extended prism exposure of 500 trials, the ultraslow system accounts for most of the net adaptation (i.e. the purple line is the most negative one of all three). Because of that system's high retention factor, the AEs are then very stable.

Using a Bayesian model with an even larger number of states to explain saccadic adaptation, Körding *et al.* (2007) argued that these internal states reflect the nervous system's current estimate of the perturbation timescale, which shifts towards slower, longer timescale estimates as a function of increasing exposure duration. In other words, optimal adaptation (learning and forgetting rates) reflects the solution of a temporal credit assignment problem: what is the most likely timescale of the perturbation that causes the observed performance error? A Bayesian solution treats longer exposure duration as greater evidence of a long-lasting perturbation. Hence, at any time point, optimal behaviour is defined as the motor output most consistent with the evidence of the estimated perturbation (magnitude and timescale). This theoretical framework has been shown to reproduce and explain a variety of saccade adaptation behaviours in both monkey and human (Körding *et al.* 2007). In summary, there may be a continuum of learning processes, operating over multiple timescales, whose contribution to behaviour varies with factors such as the source, duration, and volatility of the perturbation.

Following similar reasoning, a recent study combined state-space modelling and functional magnetic resonance imaging (fMRI) to test for neural signatures associated with a range of possible internal state learning and forgetting dynamics during adaptation to a novel visuomotor rotation (Kim *et al.* 2015). The rationale was to use state-space modelling to extract hidden internal states and identify brain regions where the blood oxygenation level dependent (BOLD) signal co-varies with these com-

putational dynamics. Using a dimensionality reduction analysis, Kim and colleagues identified four main components (i.e. group of multiple timescales), each associated with different neural networks. Notably, all these networks included either a parietal and/or a cerebellar component, suggesting that the functional activation timecourses of brain regions known to be associated with PA (see section 1.4) are compatible with the state-space modelling framework. In a different field, recent advances in biophysical models of neural networks also support the idea that different brain networks process information at different timescales (Fusi *et al.* 2005; Bernacchia *et al.* 2011; Chaudhuri *et al.* 2015). This may offer a potential path to bridge biologically plausible models of neural networks with computational descriptions of learning systems, and span explanatory levels from neurons to behaviour.

1.6 Conclusion

Prism adaptation is one of the oldest experimental paradigms used in the sensorimotor adaptation literature (von Helmholtz 1867). It warrants special attention mostly because of its rehabilitation potential for post-stroke spatial neglect (Rossetti *et al.* 1998). However, the traditional cognitive psychology theoretical framework within which PA has mainly been studied suffers from important explanatory flaws that limit mechanistic understanding. In this introductory chapter, we have advocated for the utility of a computational framework that re-conceptualises PA in terms of its constituent algorithms. We have focused on state-space models that offer a simple mathematical description of the temporal dynamics of internal estimates thought to underlie PA behaviour, allowing to move beyond the categorical 'dualist' approach of the traditional framework. This approach enables quantitative questions about information processing and neural implementation, such as: 1) how many processes

best explain PA behaviour?, 2) is there a discrete number or a continuum of processes/timescales which varies with task parameters?, or 3) are distinct brain regions associated with distinct timescales, or might a given brain circuit have the ability to perform similar computations over a range of differing timescales? We submit that progress in understanding the functional and neural bases of PA behaviour requires this experimental paradigm to be re-conceptualised at an algorithmic level of description. Doing so offers our field the opportunity to capitalise on explanatory gains generated by the literature on computational sensorimotor control. Such insights are being leveraged in the literature on other kinds of adaptation task, but the prism literature has so far remained relatively immune. Given the distinctive features of PA, and its applications in neuropsychology, we believe the time for our field to start leveraging these gains is ripe.

1.7 Aims of this thesis

The primary aim of this thesis is to provide a better understanding of the neuro-computational mechanisms underlying sensorimotor memory formation during prism adaptation in order to maximise persistence of the after-effect, which might improve clinical benefits in neglect patients. Throughout the work reported in this thesis, I aim to use a principled translational approach grounded on a better basic understanding of the neural and psychological mechanisms involved in the formation of the AE memory trace. I use a range of neuroimaging and brain stimulation techniques, as well as behavioural data and computational modelling, in the development of this thesis. First, I set out to define a principled, quantitative framework within which to investigate the PA memory trace ([chapter 4](#)). As a result, the prism after-effect will be conceptualised as a weighted combination of a short-lasting and long-lasting memory trace, quantified as a *relative weight* parameter. Next, neurochemical corre-

lates of this parameter are identified within human primary motor cortex in young healthy individuals ([chapter 5](#)). To establish causality, the effect of excitatory anodal transcranial direct current stimulation (a-tDCS) is assessed at the physiological, computational and behavioural level. This will reveal a neuro-computational mechanism by which trait and state variation in sensorimotor excitation:inhibition ratio causally determines memory persistence/decay after PA. In [chapter 6](#), the same neuro-computational mechanism is tested in a cohort of older adults. Although the mediation by the computational parameter appears to be different in this population, the positive association between primary motor cortex excitability and sensorimotor memory persistence is preserved. Finally, building on these results and on previous work from the laboratory (O'Shea *et al.* 2017), the effect of excitatory a-tDCS as an adjuvant to PA therapy is tested in a small group of 5 neglect patients. Chapter [7](#) presents the functional and structural bases of attentional deficits in these patients. In [chapter 8](#), the clinical efficacy of PA + a-tDCS is compared to PA + sham tDCS.

Chapter 2

Methods for stimulating the brain

Non-invasive brain stimulation techniques are neuromodulatory tools used to influence neuronal activity. Transcranial magnetic stimulation (TMS) is used in [Appendix B](#), and transcranial direct current stimulation (tDCS) is used in [chapter 5](#), [6](#) & [8](#). In this chapter the methods used in the aforementioned chapters are discussed, with a particular emphasis on primary motor tDCS.

2.1 Transcranial Magnetic Stimulation (TMS)

Transcranial Magnetic Stimulation (TMS) is a non-invasive neurostimulation technique involving the application of an electrically induced transient magnetic field through the scalp and skull to a targeted cortical region. It was invented by Barker *et al.* (1985), who first demonstrated that the application of TMS to the primary motor cortex (M1) could induce a twitching movement in a specific body part contralateral to the stimulated hemisphere.




Figure 2.1: Example TMS machine and application. Experimenter delivering TMS to a volunteer and recording the electromyographic response.

2.1.1 Practical applications

The basic components of a TMS set-up include a stimulator and a TMS coil (Figure 2.1). According to fundamental principles of electromagnetic induction, the brief current flow generated in the coil by the stimulator induces a transient magnetic field, typically in the order of 1-2 Tesla and lasting for about 200 μs (Hallett 2007), called a TMS pulse. If the coil is placed on the scalp of a volunteer, the TMS pulse induces a flow of electric current in the underlying brain tissue, which is responsible for a rapid depolarisation that can lead to action potentials under certain circumstances (Sandrini *et al.* 2011). Several experimental parameters such as the stimulation intensity, frequency or timing, can influence the depolarising effect of TMS (Sandrini *et al.* 2011; Valero-Cabr  *et al.* 2017). Some of these factors are discussed in the following sections.

2.1.2 Motor Evoked Potentials and resting motor threshold

When TMS is applied to the primary motor cortex with sufficient intensity, the depolarisation of motor neurons in the descending corticospinal tract can elicit a peripheral response that is measurable and quantifiable via electromyographic (EMG) recording of the muscles that receive projections from the stimulated motor neurons ([Figure 2.1](#)). The EMG activity resulting from a TMS pulse is termed a motor-evoked potential (MEP). The spatial resolution of the TMS pulse is such that individual muscles or groups of muscles can be targeted by stimulating specific areas along the somatotopically organised primary motor cortex (Penfield & Boldrey 1937). MEPs are typically recorded from the first dorsal interosseous muscle, which requires placing the TMS coil roughly over C3/4 ([Figure 2.4](#)).

Several MEP properties (e.g. delay, size, shape) can be used as markers of different aspects of corticospinal physiology. Relevant to this thesis is the size (amplitude) of MEPs. It is recognised that three basic physiological mechanisms influence the size of MEPs (Wassermann *et al.* 2008):

1. the number of recruited motor neurons in the descending corticospinal tract;
2. the number of these neurons discharging more than once in response to the TMS pulse;
3. the degree of synchronisation of the TMS-induced motor neuron discharges

Thus, when it is measured with a constant stimulation intensity, the size of MEPs provides a sensitive measure of excitability within M1, the corticospinal tract, and the connected muscles (Nitsche & Paulus 2000; Nitsche & Paulus 2001). For example, this technique has been used extensively to measure the consequences of tDCS on corticospinal excitability (see [section 2.2](#)).

The resting motor threshold (rMT) is defined as the minimal intensity required to elicit an EMG response of at least 50 μ V with 50% probability in a fully relaxed muscle (Rossini *et al.* 2015). This measure is highly reproducible within-subject (Stewart *et al.* 2001) and is influenced by the distance from coil to underlying cortex (McConnell *et al.* 2001). It is considered the simplest biomarker of individual excitability level and provides a way of calibrating stimulation intensity on an individual basis (Rossi *et al.* 2009). Typically, a TMS experiment begins with a measurement of the rMT, which is then used as a reference to determine stimulation intensity, expressed as a percentage of the rMT. Although the relationship between the rMT and the excitability of other non-motor brain regions is not obvious (Stewart *et al.* 2001; Boroojerdi *et al.* 2002; Gerwig *et al.* 2003; Antal *et al.* 2004; Deblieck *et al.* 2008) and may require taking into account the coil-to-cortex distance (Stokes *et al.* 2013), this method is used commonly, even when brain regions outside the motor cortex are targeted by TMS (Rossi *et al.* 2009; Rossini *et al.* 2015; Valero-Cabré *et al.* 2017). The main reason for the broad use of this calibration technique is that the rMT provides an easy and highly reproducible measure of excitability that is often impossible to get from regions outside motor or visual areas, where no objective, well-characterised physiological response can be registered.

2.1.3 Repetitive TMS

TMS is a polyvalent tool that can be used to induce 'virtual lesions', determine the chronometry of psychological processes, measure or manipulate cortical excitability, assess functional connectivity, etc. (Pascual-Leone *et al.* 2000; Valero-Cabré *et al.* 2017). Each of these applications is associated with a specific form of TMS (e.g. single pulse, paired pulse, high or low frequency repetitive TMS) but many are outside the scope of this thesis. In [Appendix B](#), low frequency (1 Hz) repetitive TMS (rTMS) is applied offline to transiently inhibit function of the right posterior parietal cortex.

We therefore focus here on this type of stimulation only.




Figure 2.2: Effect of repetitive TMS on MEPs. MEPs obtained before and after the delivery of low frequency (1 Hz) rTMS to the primary motor cortex at 90% of the rMT for 10 minutes (600 pulses). Note the decrease in the amplitude of the MEPs, reflecting decreased corticospinal excitability. mV: millivolts; ms: milliseconds. Figure adapted from Valero-Cabré *et al.* (2017).

Repetitive TMS refers to a type of stimulation in which TMS pulses are delivered by train, and, as a result, have the ability to produce different effects than isolated pulses. Many parameters can be manipulated (e.g. frequency, duration, number of pulses, number and duration of intervals between pulse trains), leading to a large number of potential configurations. When applied continuously at low frequency (1 Hz) for 5 minutes or more, rTMS has been shown to produce a sustained reduction of excitability within the targeted cortical area, either weakening its intrinsic level of activity, reducing the efficacy of its outputs, or silencing its contribution to cognitive processes in which it is involved (Chen *et al.* 1997; Boroojerdi *et al.* 2000; Maeda *et al.* 2000; Fitzgerald *et al.* 2002; Lang *et al.* 2006; Valero-Cabré *et al.* 2007; Petit *et al.* 2015) (example in the motor domain in Figure 2.2). Typically, cortical activity remains altered for a period of at least 30 minutes after the end of the stimulation, this is termed rTMS after-effects (Valero-Cabré *et al.* 2017). In *offline* rTMS experiments, behavioural performance is compared before and after the administration of rTMS at rest, to evaluate the contribution of a targeted area to a task of interest (for example, see Kosslyn *et al.* 1999; Hilgetag *et al.* 2001). This type of experimental paradigm is particularly well suited for functional magnetic resonance imaging (fMRI) studies because rTMS can be administered outside the scanner (Sack *et al.* 2002b; Plow *et al.*

2014; Petitet *et al.* 2015). The drawback of this approach compared to combining fMRI with single pulse TMS inside the scanner (Formisano *et al.* 2002; Sack *et al.* 2002a; Sack *et al.* 2005) is the relatively weak effect it has on brain function, in particular due to the passage of time and the development of compensatory mechanisms (Plow *et al.* 2014). However, its major advantage is that it can be applied outside the MRI scanner, avoiding all the technical complications and potential artefacts related to generating electromagnetic pulses in the MR environment.

2.1.4 Safety considerations

TMS is a safe technique when it is applied within the recommended international safety guidelines (Wassermann 1998; Rossi *et al.* 2009). These involve screening volunteers thoroughly for eventual risk factors when recruiting them for any brain stimulation study (Rossi *et al.* 2011), and taking precautions to minimise the risk of inducing side effects during experimental testing. Here, we summarise the main potential side effects associated with rTMS (Rossi *et al.* 2009), and how to safely manage them.

- **Seizure.** The most severe safety concern with rTMS is the induction of seizures due to hypersynchronised discharges of groups of neurons caused by increased excitation:inhibition ratio (Wassermann *et al.* 2008). Seizures can be induced by high frequency (excitatory) rTMS 1) during the stimulation, or 2) when the after-effects are present. Although the first has been reported (predominantly when rTMS was delivered outside the 1998 safety guidelines; Wassermann 1998; Rossi *et al.* 2009), there is no record of the latter ever occurring. Additionally, low frequency is often considered to have a protective effect on seizures and is associated with a very low risk of kindling (Theodore *et al.* 2002). Screening volunteers for personal or family history of seizures is a way of reducing the probability of this side effect to occur. Care should also be taken with

medications that alter depolarisation thresholds, such as anti-depressants and anti-epileptics, and in particular drugs reducing GABAergic inhibition (Dobek *et al.* 2015). Healthy volunteers taking such medications should be excluded from rTMS studies. In case of favourable risk/benefit ratio, exceptions can be made for patients, with their informed consent, and in consultation with a medical doctor. In the unlikely event of a seizure, TMS should be terminated immediately. The participant should be assisted in controlled reclining, with sufficient space around him/her to prevent potential injuries caused by knocking objects. Airway breathing and circulation should be assessed. It is also important to monitor the duration of the seizure and inform the emergency service. As soon as movements (convulsions) cease, subject should be turned on one side to help clear the airway and avoid aspiration.

- **Syncope.** Neurocardiogenic syncope is a common reaction to anxiety and psycho-physical discomfort, and it occurs more often than seizures during TMS experiments (Rossi *et al.* 2009). The clinical feature that distinguishes syncope from seizure is rapid recovery of consciousness within seconds (and not minute(s)). Typically, the prodromes include verbal complaints (e.g. “I need to lie down” or “I need air”), followed by a narrowing and blacking out of the visual field, sensations of heat, bradycardia, and transient loss of peripheral pulse. Initial management of suspected seizure and syncope are identical. Recovery of normal consciousness beyond 30 seconds following onset mandates further medical evaluation. The risk of inducing a syncope can be minimised by explaining carefully the experimental procedure to participants, and regularly asking them how they feel.
- **Local pain, headache, discomfort.** Single-pulse TMS is generally well tolerated and experienced as painless by most participants. However, scalp or neck

pain, headache and discomfort are by far the most common side effects of rTMS (Rossi *et al.* 2009). It is unclear what causes the painfulness of TMS, and it depends on individual susceptibility, stimulation location, coil design, intensity and frequency of stimulation. Trigeminal stimulation, acoustic noise, muscle contractions, forced posture and head immobilisation are likely to play a role in the painful experience of rTMS. Most of the time, these side effects vanish rapidly after offset of the stimulation. Headache may occasionally persist and, in this case, a common analgesic administered orally is recommended.

- **Hearing loss.** Rapid mechanical deformation of the TMS coil when a pulse is generated produces a short but intense acoustic artefact that can exceed 140 decibels (dB) (Dhamne *et al.* 2014). For reference, 120 dB is the noise level of a thunderclap and 140 dB is reached in rock concerts. In the absence of hearing protection, this level of impulsive noise exceeds the safety levels for the auditory system recommended by the Occupational Safety and Health Administration (OSHA). The use of approved hearing protection (earplugs or ear muffs) is therefore recommended for all participants. Individuals with known pre-existing noise induced hearing loss, or undergoing ototoxic medications (e.g. aminoglycosides, cisplatin), or with cochlear implants should be excluded from TMS studies. A small number of transient increases in auditory thresholds have been reported following TMS (Rossi *et al.* 2009). In such cases, a prompt referral for auditory assessment is recommended.
- **Cognitive changes.** In the majority of brain stimulation experiments, the purpose of applying rTMS to the brain is to induce a change in the cognitive domain. These changes are generally small and constrained within the frame of the experiment (i.e. short-lasting). However, theoretically rTMS is associated with a risk of inducing un-expected, potentially long-lasting, cognitive changes.

This type of side effect is extremely rare (Rossi *et al.* 2009) but should be monitored carefully.

- **Others.** A final safety concern is the risk of inducing heating, attractive (or repulsive) forces or voltage in any implanted device (e.g. electrodes, skull plate, aneurysm clips, cochlear implants, deep brain stimulation implant, wires, heart pacemaker, etc.) situated nearby the TMS coil. All potential participants should therefore be screened for such devices and excluded if they represent a safety concern.

In this thesis, TMS was delivered in appliance to the international safety guidelines (Rossi *et al.* 2009). All participants were young healthy controls and no side effect was observed.

2.2 Transcranial Direct Current Stimulation (tDCS)

Transcranial direct current stimulation (tDCS) is a non-invasive neuromodulatory technique involving the application of a weak static direct electrical current through the brain via two or more electrodes, typically placed on the scalp (Figure 2.3).




Figure 2.3: Example tDCS kit and application. A. Battery operated direct current stimulator connected to an anode and cathode electrode. B. Experimenter delivering tDCS to a volunteer.

The first documented evidence of the use of transcranial electrical stimulation comes in Roman Empire times, when the physician Scribonius Largus described

how placing a living torpedo fish over the scalp could relieve headache in a patient (Sarmiento *et al.* 2016). Sporadic reports of the use of electric fishes to treat neurological conditions have been found within the next centuries. At the end of the 18th century, Luigi Galvani invented a Direct Current (DC) battery (galvanic battery) and his nephew, Giovanni Aldini, was presumably one of the first persons to use it to treat neurological and psychiatric conditions (Fitzgerald 2014). In the next 130 years, 'electrotherapy' became relatively popular, particularly among German psychiatrists, but the variation of procedures, unclear descriptions and few qualitative details yielded rather controversial and un-reproducible results (Sarmiento *et al.* 2016). The clinical use of DC stimulation was therefore abandoned from the 1930's (Steinberg 2013). The research community briefly regained some interest in weak DC stimulation in the 1960's, mainly due to electrophysiological studies in animals reporting long-lasting changes in neuronal excitability (Bindman *et al.* 1962; Bindman *et al.* 1964). However, the development of neuropsychopharmacology in the 1970's led once again to an abandon of the field (Priori 2003; Dubljević *et al.* 2014). It was not until 1998 that Italian and German research teams resurrected the field from obscurity (Priori *et al.* 1998; Nitsche & Paulus 2000), and modern tDCS as we will discuss it in this section was born. The number of studies using this technique has since risen exponentially (Dubljević *et al.* 2014).

2.2.1 Practical applications

The basic components of a tDCS set-up include a battery operated DC simulator and two or more electrodes (Figure 2.3). The current generated by the stimulator is delivered via relatively large conductive rubber electrodes (typically sized between 25 and 35 cm²), inserted in saline soaked sponges or covered in electrode gel, and held on the scalp. Saline soaked sponges or electrode gel are used to reduce electrical resistance, hence minimising current dispersion and participant discomfort. Electrode gel

is preferable over saline soaked sponges when tDCS is delivered inside the MRI scanner (e.g. [chapter 5](#)), as the latter would dry out rapidly over the scanning period. By convention, a positively charged electrode is called anode and a negatively charged one is called cathode. The electrode montage (number, shape and position of the electrodes) is a key experimental parameter influencing the distribution and polarity of the induced electric field (Datta *et al.* 2009) as well as the resulting physiological (Nitsche & Paulus 2000; Stagg *et al.* 2009) and behavioural effects (Galea *et al.* 2010; Kang & Paik 2011; Stagg *et al.* 2011d; O’Shea *et al.* 2017).




Figure 2.4: Left M1 anodal tDCS montage. According to the international EEG coordinate system, the anode (positive) electrode was placed over C3 and the cathode (negative) electrode over FP2.

Throughout this thesis, we used a standard left primary motor cortex (M1) tDCS protocol (Chapters 5, 6 & 8). The montage consisted in an anode electrode placed over the left M1 (C3 in the 10-20 international EEG system) and a cathode electrode placed over the contralateral (right) supraorbital bridge (for similar tDCS parameters, see Nitsche & Paulus 2000; Galea *et al.* 2010; Stagg *et al.* 2011d; O’Shea *et al.* 2014; O’Shea *et al.* 2017) ([Figure 2.4](#)). The correspondence between C3 and M1 (Yousry *et al.* 1997) has been confirmed by neuroimaging studies (Herwig *et al.* 2003; Okamoto *et al.* 2004). The stimulation was delivered at 1 mA for 20 minutes with

a current ramp up and down duration of 10 seconds at the beginning and end of the stimulation, respectively. The gradual current increase and decrease help prevent electrical transients that can result in noticeable cutaneous sensations (itching or tingling) and retinal phosphenes (Poreisz *et al.* 2007; Woods *et al.* 2016; Zhao *et al.* 2017). During sham stimulation, the DC stimulator is typically switched off after the initial ramp up period and no direct current is delivered. Instead, small current pulses (110 μ A over 15 ms) can be delivered every 550 ms to simulate the mild tingling sensation that can accompany real stimulation (Gandiga *et al.* 2006; Kessler *et al.* 2012). Typically, volunteers cannot distinguish between real and sham stimulation.

2.2.2 Physiological basis of tDCS

Unlike other non-invasive brain stimulation techniques, such as TMS, the weak static electric field produced by tDCS cannot generate the rapid depolarisation required to induce action potentials (Rahman *et al.* 2017). Instead, the current view is that tDCS modulates neuronal excitability via a direct action (hyperpolarisation or depolarisation) on resting membrane potential (Nitsche & Paulus 2000; Nitsche *et al.* 2003; Radman *et al.* 2009; Rahman *et al.* 2013; Rahman *et al.* 2017), and a long-term effect on synaptic plasticity (Nitsche *et al.* 2004b; Fritsch *et al.* 2010; Stagg & Nitsche 2011; Krause *et al.* 2013). Here, we summarise the body of work that investigated the physiological basis of tDCS. A particular emphasis is given to M1 anodal tDCS because it is the only montage used in this thesis (Chapters 5, 6 & 8). The following three sections are not organised chronologically but rather review levels of investigation from the molecular/cellular scale to the whole-brain level.

2.2.2.1 Electrophysiological effect of Direct Current stimulation (DCS) *in vitro*

Early electrophysiological studies reported spontaneous and evoked cortical potentials to be acutely facilitated by positive current (anodal stimulation) and inhibited by negative current (cathodal stimulation) (Bindman *et al.* 1962; Creutzfeldt *et al.* 1962; Bindman *et al.* 1964; Purpura & McMurtry 1965). In other words, for a given presynaptic input, the amplitude of the excitatory postsynaptic potential (EPSP) is increased by depolarising current (anodal DCS) and decreased by hyperpolarising current (cathodal DCS). Although not fully understood at the time, these effects were found to outlast the duration of the DC stimulation (Bindman *et al.* 1964; Gartside 1968), suggesting an effect at the synaptic level. In a series of electrophysiological experiments in mouse M1 slices, Fritsch *et al.* (2010) investigated the cellular and molecular mechanisms underlying such effects. The application of anodal DCS (electric field ≈ 0.75 V/m) concurrently to a 0.1 Hz Low Frequency Stimulation (LFS) resulted in potentiation of the amplitude of field excitatory postsynaptic potentials (fEPSP), an electrophysiological marker of synaptic efficacy, that outlasted the duration of the stimulation for up to 2 hours (Figure 2.5). This effect was weaker when a different LFS frequency (0.01 Hz or 0.2 Hz) was used and was not found with anodal DCS alone (no LFS), indicating that the underlying neuronal activity occurring during DCS constitutes a key mediator of the synaptic strengthening effect. This is arguably analogous to combining tDCS *in vivo* with a cognitive task or learning paradigm which functionally activates a targeted network. Therefore, for a given montage, the spatial distribution of the synaptic effect of tDCS is likely to depend upon the pattern of endogenous activity induced by the task. Cognitive state dependency of tDCS effects (also well established for TMS effects) will be discussed more extensively elsewhere in this thesis (chapter 5).




Figure 2.5: DCS promotes LTP in motor cortical slices. **A.** Schematic diagram of DCS and recording from mouse M1 slices. The distribution of the electric field strength (0.75 V/m in the motor cortex) is superimposed onto a coronal slice of M1. **B.** Normalised field excitatory postsynaptic potentials (fEPSP) amplitude in M1 layer II/III before, during 15 min of anodal DCS coupled with 0.1 Hz Low Frequency Stimulation (LFS), and up to 30 min after DCS. The stimulation resulted in potentiation of the fEPSP, which outlasted the stimulation duration. This effect was not observed in the absence of LFS and was abolished by pre-treatment with the NMDA receptor antagonist D-APV (Figure 2.6). Figure adapted from Fritsch *et al.* (2010).

Two molecular factors, known to be important for synaptic plasticity induction, have been identified as key mediators of the DCS effect *in vitro*: N-methyl-D-aspartate (NMDA) and Brain-derived Neurotrophic Factor (BDNF) (Minichiello 2009). The binding of glutamate to the postsynaptic NMDA receptor induces a postsynaptic calcium influx, which is necessary for the induction of Long-Term Potentiation (LTP) (Bliss & Collingridge 1993). This cascade of event (NMDA activation and postsynaptic calcium influx) also induces the release of endogenous BDNF (Balkowiec & Katz 2002; Tanaka *et al.* 2008), a neurotrophic factor involved at various stages of LTP induction, including its maintenance (for review, see Bramham & Messaoudi 2005; Minichiello 2009). In Fritsch and colleagues' experiments, pre-treatment of M1 slices with a NMDA receptor antagonist (D-APV) prevented potentiation induction (Figure 2.6A), suggesting a role of NMDA-dependent synaptic plasticity in the induction of a sustained synaptic facilitation (Fritsch *et al.* 2010). The blockage of the BDNF-TrkB pathway in genetically modified mice also abolished the induction of LTP by DCS (Figure 2.6B). Taken together, Fritsch and colleagues' electrophysiological work

demonstrated an effect of DC stimulation on BDNF-dependent synaptic plasticity, which is a form of NMDA-dependent LTP mediated by the activity-dependent release of BDNF (Fritsch *et al.* 2010).




Figure 2.6: DCS-LTP requires NMDA and BDNF. **A.** Normalised field excitatory postsynaptic potentials (fEPSP) amplitude in M1 layer II/III of M1 slices pre-treated with the NMDA receptor antagonist D-APV before, during 15 min of anodal DCS coupled with 0.1 Hz Low Frequency Stimulation (LFS), and up to 30 min after DCS. The blockage of NMDA receptors abolished potentiation of the fEPSP, demonstrating the involvement of NMDA-dependent LTP in this effect. **B.** DCS-LTP in M1 slices from $\text{TrkB}^{\text{F616A}}$ mice in which the BDNF-TrkB pathway is blocked by the application of 1NMPP1. $\text{TrkB}^{\text{F616A}}$ slices exhibited normal DCS-LTP (gray filled diamonds). Continuous treatment with 1NMPP1 (5 mM) completely prevented DCS-LTP (black filled diamonds). 1NMPP1 (5 mM) applied after DCS prevented the post-DCS increase in fEPSP amplitude (open diamonds). Figure adapted from Fritsch *et al.* (2010).

Recent combined modelling-electrophysiological work shed light on how the diffuse and weak profile of current flow induced by DCS (Datta *et al.* 2009) could result in such an LTP-like increase in synaptic efficiency (Rahman *et al.* 2013; Rahman *et al.* 2017). In brief, these results suggest that the polarisation of afferent axons and/or postsynaptic neurons increases presynaptic and/or postsynaptic firing rate when neurons are active. This results in increased coincidence probability and hence promotes LTP naturally occurring between connected and co-active brain regions (Rahman *et al.* 2017). Importantly, the distribution of the DCS LTP-like effect can only be considered in the context of a specific pattern of endogenous activity, which is a pre-requisite to the induction of synaptic plasticity.

2.2.2.2 Mapping the current inside the head

The *in vitro* work mentioned in the previous section involved the application of DC to neural tissue in slice preparations so the experimenter could measure directly and control the resulting current density (Figure 2.5A). Although this body of work expanded our understanding of the effect of DC on synapses and neurons, the direct extrapolation of these findings to humans is limited by many methodological differences. In particular, an important and controversial question in the field of non-invasive brain stimulation is to determine the intracranial magnitude and distribution of the induced electric field when current is delivered through the scalp and skull. This debate reached a climax at the 2016 annual meeting of the Cognitive Neuroscience Society when György Buzsáki presented his un-published 'cadaver study' in which he and his colleagues applied transcranial electrical stimulation to human cadaver heads while recording inside the brain, and reported a loss of up to 90% of the current due to a shunting effect of the skin (Underwood 2016). However, dead tissue is likely to conduct electricity differently from living one, challenging the validity of human cadaver heads as a model of real-life *in vivo* tDCS.




Figure 2.7: Simulation of the spatial distribution of tDCS-induced intracranial electric field. Spatial distribution of the electric field modelled with a MRI-derived gyri-precise head model (Datta *et al.* 2009). Figure adapted from Rahman *et al.* (2013).

Early spherical head models showed that the intracranial electric field generated

by tDCS is low (i.e. below 1 V/m) with conventional tDCS (Miranda *et al.* 2006). This was confirmed by computational models using MRI-based structural details to predict electric fields inside the brain (Datta *et al.* 2009; Parazzini *et al.* 2011; Wagner *et al.* 2013; Thielscher *et al.* 2015), and by a recent study measuring electric fields inside the *in vivo* human brain during tDCS (Huang *et al.* 2017). Note that this order of magnitude of intracranial electric fields is within the lower range of what is required to affect neuronal activity in animal experiments (Fritsch *et al.* 2010; Reato *et al.* 2010).

Another important feature of the induced electric field is its heterogeneity and diffusivity. Intracerebral recordings and computational models taking into account geometry and conductivity properties of the scalp, skull, cerebrospinal fluid and white and grey matter showed diffuse distribution of the current, with discrete clusters of electric field intensity maxima (Datta *et al.* 2009; Datta *et al.* 2011; Rahman *et al.* 2013; Huang & Ahmed 2014; Thielscher *et al.* 2015) (Figure 2.7). Remarkably, simulations predict that 1) the maximum electric field intensity is not necessarily found underneath the stimulated electrode as commonly assumed (Datta *et al.* 2009), and 2) different sulci might experience depolarising or hyperpolarising effects depending on their orientation relative to the electric field (Rahman *et al.* 2013; Rahman *et al.* 2015; Rahman *et al.* 2017). These observations encouraged certain research groups to develop more focused stimulation protocols using multiple smaller electrodes, a technique commonly known as high-definition tDCS (Datta *et al.* 2008; Dmochowski *et al.* 2011; Edwards *et al.* 2013; Kuo *et al.* 2013).

2.2.2.3 Electrophysiological effect of tDCS *in vivo*

Most electrophysiological studies of tDCS in humans used the primary motor cortex as a model system because of the direct read-out it offers on cortical excitability via

the measurement of MEPs, i.e. the electromyographic activity resulting from a M1 TMS pulse (see section 2.1.2). Typically, these experiments involve the measurement of MEPs, before, during and after the application of M1 tDCS. Increases or decreases in MEPs are interpreted as markers of increased or decreased M1 excitability, respectively.




Figure 2.8: Effect of tDCS on cortical excitability in humans. A. Size of MEPs during anodal/cathodal tDCS normalised to no stimulation. MEP amplitudes increased with anodal and decreased with cathodal tDCS. Asterisks indicate significant differences between values with and without stimulation ($p < .05$). **B.** Time course of polarity-specific changes in M1 excitability. MEPs in the post-tDCS period are normalised to baseline. Following anodal (squares) and cathodal (circles) tDCS, MEPs returned to baseline in about 5 minutes. Figure adapted from Nitsche & Paulus (2000)

Consistent with the *in vitro* results reported earlier (Bindman *et al.* 1962; Creutzfeldt *et al.* 1962), M1 anodal tDCS has been found to reliably induce an increase in MEPs (reflecting enhanced cortical excitability) that outlasts the duration of the stimulation (Priori *et al.* 1998; Nitsche & Paulus 2000; Nitsche & Paulus 2001; Liebetanz *et al.* 2002; Nitsche *et al.* 2004b; Nitsche *et al.* 2005; Monte-Silva *et al.* 2013; O'Shea *et al.* 2014) (Figure 2.8). When the electrical current is reversed (cathodal tDCS), the opposite effect (i.e. decrease in MEPs) is found (Nitsche & Paulus 2000; Nitsche & Paulus 2001; Liebetanz *et al.* 2002; Nitsche *et al.* 2003; Nitsche *et al.* 2005; O'Shea *et al.* 2014). Pharmacological manipulations using

antagonisation approaches probed these effects further to determine the ion channels and receptors involved in humans. Consistent with the *in vitro* results (see section 2.2.2.1), the after-effect of tDCS was found to be NMDA-dependent, while the online effect was found to implicate sodium channels involved in the regulation of membrane potential (Liebetanz *et al.* 2002; Nitsche *et al.* 2003). Pharmacological strengthening of GABAergic inhibition via administration of the GABA_A receptor agonist Lorazepam abolished the online excitatory effect of anodal tDCS (Nitsche *et al.* 2004b), suggesting a potential role of GABA in mediating the effect of anodal tDCS in humans.




Figure 2.9: Change in GABA before, during, and after anodal tDCS. Graph shows the change in left M1 GABA:Creatine ratio relative to the two baseline measurements during and after either real (black circles) or sham (white circles) left M1 anodal tDCS (20 minutes, 1 mA). The timecourse of change in GABA relative to baseline shows a gradual decrease during anodal stimulation that became significant 10–15 min after stimulation finished (Post1). Figure adapted from Bachtiar *et al.* (2015)

Since then, several magnetic resonance spectroscopy studies reported anodal tDCS to decrease GABA concentration underneath the stimulated electrode (Stagg *et al.* 2009; Kim *et al.* 2014; Bachtiar *et al.* 2015; Barron *et al.* 2016; Antonenko *et al.* 2017) (Figure 2.9).

2.2.3 Effect of M1 anodal tDCS on sensorimotor adaptation

M1 anodal tDCS has been used mainly to modulate motor learning, i.e. the ability to acquire new skills or refine motor actions through repeated practice (Reis *et al.* 2008; Ammann *et al.* 2016). This is not the primary focus of this thesis. More relevant to our work, studies that used this technique to modulate sensorimotor adaptation are summarised in [Table 2.1](#). In most of these studies, M1 anodal tDCS did not influence the rate of adaptation but rather influenced the retention/generalisation pattern of the subsequent after-effects (Hunter *et al.* 2009; Galea *et al.* 2010; Orban de Xivry *et al.* 2011; O'Shea *et al.* 2017).

Remarkably, in O'Shea *et al.* (2017), the magnitude of stimulation-induced reduction in M1 GABA concentration correlated with the increase in retention 24 hours after the adaptation, implicating GABAergic inhibition in the generation of this behavioural effect. This finding will be characterised further in [chapter 5](#), using computational modelling and multimodal 7 Tesla magnetic resonance imaging.

Study	Paradigm	Arm	Site	Param.	Exp. design	Result
Hunter <i>et al.</i> 2009	FF	R	L	1 mA 35 cm ² 17 min	<i>Within-subject</i> n = 14 A/S	Greater overshoot during early stage of de-adaptation
Galea <i>et al.</i> 2010	VMR	R	L	2 mA 25 cm ² 17 min	<i>Between-subject</i> n = 10 S n = 10 A	Increased AE (without visual fbk)
Orban de Xivry <i>et al.</i> 2011	FF	R	L	1 mA 25 cm ² 20 min	<i>Between-subject</i> n = 11 S n = 8 A	Increased generalisation in intrinsic (but not extrinsic) coordinates
Block & Celnik 2013	VMR	L/R	L/R	2 mA 25 cm ² 15 min	<i>Between-subject</i> Exp 1: n = 15 S; n = 15 A Exp 2: n = 6 S; n = 6 A	No effect
Herzfeld <i>et al.</i> 2014	FF	R	L	2 mA 25 cm ² 25 min	<i>Between-subject</i> n = 12 S n = 14 A	No effect
O'Shea <i>et al.</i> 2017	PA	R	L	1 mA 35 cm ² 20 min	<i>Within-individual</i> Exp 1: n = 9 S/A Exp 2: n = 10 S/A	Increased retention of AE (without visual fbk) in both experiments

Table 2.1: Summary of studies investigating the effect of primary motor cortex anodal tDCS during sensorimotor adaptation. A: anodal tDCS, AE: after-effect, fbk: feedback, FF: force field adaptation, L: left, PA: prism adaptation, R: right, S: sham tDCS, VMR: visuomotor rotation adaptation,

2.2.4 Safety and ethical considerations

Unlike TMS, there is no evidence that tDCS can induce seizures, even in cases of compromised neurophysiology. To date, no serious adverse effect or irreversible injury has ever been reported with conventional tDCS protocols (≤ 40 min, ≤ 4 mA, ≤ 7.2 Coulombs) in the general population, potentially susceptible individuals (e.g. subjects with altered neuroanatomy or neurophysiology), or susceptible individuals (e.g. children) (Bikson *et al.* 2016). Nevertheless, the same exclusion criteria and safety guidelines as TMS apply to tDCS regarding susceptibility to seizure and syncope (Woods *et al.* 2016) (see section 2.1.4). Additionally, because of the risk of inducing heating and voltage, participants with implanted conductive devices should be excluded.

The most commonly reported subject discomfort is tingling under one or both electrodes (Poreisz *et al.* 2007; Woods *et al.* 2016; Zhao *et al.* 2017), presumably due to direct electrical stimulation of the cutaneous nerves. Mild skin erythema is also relatively common during tDCS but is not directly hazardous and resolves spontaneously after stimulation (Guarienti *et al.* 2015). The application of tDCS does not generate acoustic noise and is therefore not associated with a risk of hearing damage. Extended safety studies in humans performed within the recommended guidelines (Woods *et al.* 2016) found no evidence of 1) heating under the electrodes that could cause burning (Nitsche & Paulus 2000); 2) oedema or injurious alterations of the blood–brain barrier or cerebral tissue (Nitsche *et al.* 2004a); or 3) neuronal damage (Nitsche & Paulus 2001). tDCS is therefore regarded as a simple, safe, painless and inexpensive neuromodulatory technique that makes it an attractive tool in a controlled clinical or research setting. However, this great accessibility is accompanied by several ethical considerations and risks (Cohen Kadosh *et al.* 2012). Since its development in the early 2000’s, many blog

articles and videos on how to build 'do it yourself' tDCS kits have flourished on the Internet (e.g. <http://www.diytdcs.com>, <http://www.instructables.com/id/Build-a-Human-Enhancement-Device-Basic-tDCS-Suppl/>). These posts are often accompanied by unrealistic expectations of exceptional cognitive improvements or sporting achievements. The difficulty to regulate personal (mis)use of such devices at home outside the recommended guidelines exposes individuals to potential risks, and contributes to a certain misperception of tDCS. Finally, neuroenhancement techniques such as tDCS raise multiple more general ethical questions (Erhardt & Štrac 2016). Issues of implicit and explicit coercion in a society with growing enhancement practice where 'better than the natural normal' can easily become the new standard are important ethical questions that need to be addressed.

Chapter 3

Methods for imaging the brain

This chapter describes the main magnetic resonance techniques used in this thesis, emphasising their advantages and limitations.

3.1 Physics principles of Magnetic Resonance Imaging

Imaging the brain's detailed anatomy, recording its activity and measuring its neurochemical composition are key tools available to neuroscientists to investigate brain-behaviour relationships. Magnetic Resonance Imaging (MRI) is a widely used technique for this purpose and it relies mainly on imaging of the proton, the positively charged spinning nucleus of hydrogen atoms. Hydrogen is an ideal isotope for imaging the brain because of its abundance in organic tissues containing water, lipids and macromolecules. In the following sections we will describe the basic physics principles underlying the generation and detection of MR signal.

3.1.1 Polarisation and precession

MRI techniques rely on two main properties of protons: their charge and their spin. Because of these two properties, protons exhibit a magnetic moment that is typically represented as a directional vector (Figure 3.1). In the absence of external magnetic field, Brownian motion causes protons to spin in random directions, resulting in no net magnetic moment (noted M_0). However, when exposed to an external magnetic field (noted B_0), protons become polarised, i.e. their intrinsic spins align parallel or anti-parallel to the direction of the external magnetic field (Figure 3.1B). The ratio of parallel versus anti-parallel spins determines the net magnetic moment M_0 . Increasing the magnetic field B_0 results in a greater number of spins aligned parallel, and, therefore, larger net magnetic moment M_0 . In addition to aligning to an external magnetic field, protons also exhibit resonance behaviour, i.e. they precess around their preferred direction. The Larmor equation describes this magnetic resonance effect as follow:

$$\omega_0 = \gamma \times B_0 \quad (3.1)$$

where ω_0 is the resonance frequency in megahertz (MHz), γ is the gyromagnetic ratio in megahertz per Tesla (MHz/T) and B_0 is the strength of the external magnetic field in Tesla (T). In other words, there is a linear relationship between the resonance frequency of a given nucleus (ω_0) and the strength of the magnetic field it is experiencing (B_0). Because atomic nuclei differ in their neutron and proton composition, the gyromagnetic ratio γ is a nucleus-specific constant and it equals to 42.57 MHz/T for Hydrogen.




Figure 3.1: Magnetic moments of protons. **A.** In the absence of external magnetic field B_0 , spins are randomly oriented, resulting in no net magnetisation M_0 . **B.** If an external magnetic field B_0 is applied, spins align parallel or anti-parallel to its main direction. Because they cancel each other, the ratio of parallel to anti-parallel spins defines the net magnetisation M_0 .

3.1.2 RF pulses and excitation

In MR physics, excitation refers to the process of tipping the net magnetic moment M_0 out of its alignment with the magnetic field B_0 using a radio-frequency (RF) pulse. The RF pulse results in an electrically-induced magnetic field, noted B_1 , oscillating in a plane perpendicular to the constant magnetic field B_0 . By convention, B_0 runs along the longitudinal z-axis and B_1 is in the transverse (x-y) plane. Crucially, the RF pulse and the resulting B_1 only excite nuclei that precess at the precise RF pulse frequency. Thus, by matching the RF pulse frequency to the resonance frequency of interest ω_0 , one can selectively excite certain nuclei while not affecting others. Because of the linear relationship between ω_0 and B_0 (see [Equation 3.1](#)), the resonance frequency can be inferred from the strength of the magnetic field B_0 if the gyromagnetic ratio γ is known. This principle is used in MRI to select in space which slice to excite: gradients are used to vary the magnitude of B_0 along one dimension, thus inducing protons to precess at a different frequency along this dimension. A given RF pulse will therefore only excite protons in the slice that precesses at the same frequency. Most RF pulses tip the net magnetisation M_0 from 0° to 90° , and this parameter, called the flip angle, can be modified by altering the precise sequence of the RF pulse

(see section 3.1.5).

3.1.3 Relaxation and signal generation

After having been excited, protons naturally settle back along the z-axis if the external magnetic field B_0 is still present. This process is called relaxation. Because of their precession behaviour, relaxing protons follow a spiral cone shaped trajectory (Figure 3.2). The amount of net magnetisation along the z-axis gradually recovers while the magnetisation in the (x-y) transverse plane gradually decays. The rate at which these processes occur can therefore be described by two relaxation time constants, T1 and T2, that apply to spin-lattice relaxation (longitudinal recovery) and spin-spin relaxation (transverse decay), respectively. In tissue, T2 relaxation always happens faster than T1 relaxation and the factors influencing their respective time courses appear to be independent. Tissue types (e.g. white matter, grey matter, cerebrospinal fluid) differ by their T1 and T2 relaxation, a property that can be used to generate image contrast (see section 3.1.5).

The main reason for polarising, exciting and relaxing protons in MR physics is to generate a measurable electrical signal by virtue of electromagnetic induction. According to Faraday's law, if a static conductive coil is placed in a moving magnetic field (M_1 during relaxation), a current is produced in that coil. Therefore, by positioning the receiving coil in the (x-y) transverse plane, one can measure an oscillating electrical current as the transverse magnetisation decays back to zero (Figure 3.2). The intensity of this signal is determined by the amount of net transverse magnetisation M_{xy} and its frequency is dictated by the Larmor equation (Equation 3.1).

In physiological tissue, the magnetic field B_0 is not perfectly uniform and protons do not all precess at the exact same resonance frequency once excited, but rather pre-




Figure 3.2: Free induction decay. Spins follow a spiral cone shaped trajectory during relaxation. This trajectory is characterised by a T1 and T2 component, describing longitudinal recovery and transverse decay respectively. By placing a fixed receiving coil in the (x-y) transverse plane, one can measure an oscillating decaying electrical signal whose envelope provides a direct measure of transverse relaxation.

cess within a *range* of resonance frequencies. This phenomenon induces a dephasing of the excited spins, which results in a loss of net transverse magnetisation. Because it artificially increases transverse relaxation, this phenomenon is called T2*, despite happening completely independently from T2 relaxation. It is responsible for image dropouts, i.e. a loss a MR signal predominantly affecting the air-tissue boundaries where the magnetic field is the most inhomogeneous. T2* can however be prevented by applying a second 180° refocusing' RF pulse after the initial 90° RF pulse. This type of MR sequence is called 'Spin Echo' as opposed to 'Gradient Echo' when no refocussing pulse is used. By reversing the dephasing, the refocusing pulse induces a regain of magnetisation in the transverse plane. We will describe in sections 3.2, 3.3, and 3.4 the importance of dephasing for measuring water diffusion, local blood oxygenation and metabolite concentration in the brain.

3.1.4 Read-out methods

The 3-dimensional localisation of the source of the MR signal relies on the linear relationship between B_0 and resonance frequency (Equation 3.1). Typically, a MRI scanner contains gradients that can modify the magnitude of the magnetic field B_0 along the x , y and z dimension, which results in changes in resonance frequency along these dimensions. Slice selection is achieved by applying a gradient along the z dimension during the RF pulse: only the slice that resonates at the same frequency as the RF pulse will be excited (Figure 3.3). The thickness of the slice (i.e. voxel size in the z dimension) depends on the range of frequencies included in the RF pulse.




Figure 3.3: Principles of spatial encoding. Slice selection is achieved by varying precession frequency along one dimension (here z) during the RF pulse. In-slice spatial encoding is achieved by modulating the precession phase and frequency along the remaining 2 dimensions (here x and y). As a result, signal originating from voxels along the 'frequency-encode' and 'phase-encode' dimension would have the same frequency or phase respectively.

Within the selected slice, the x and y dimensions are encoded by the phase encoding gradient (GPE) and the frequency encoding gradient (GFE). The GPE intervenes

transiently right after the RF pulse (and before signal acquisition) in order to induce a dephasing along a certain dimension (y -dimension in [Figure 3.3](#)). The GFE is applied continuously during signal acquisition and modifies the precession frequency along the remaining dimension. Decomposing the signal into its constituent phases and frequencies allows the localisation of the source of signal and thereby the generation of an image where colours represent signal intensity ([Figure 3.3](#)). This type of signal read-out is well suited for high resolution structural imaging but comes at the cost of being relatively slow. Because only one line of a slice is acquired after each excitation RF pulse, it takes several minutes to acquire an image of the entire brain, called a volume.

Under certain circumstances, temporal resolution needs to be shortened in order to measure phenomena occurring on a fast timescale, while individuals are lying in the scanner (for example, see sections [3.2](#) and [3.3](#)). Under these time constraints, a common way to accelerate data acquisition is to use an Echo-planar Imaging (EPI) read-out method. In brief, the precise temporal activation sequence of GPE and GFE of EPI sequences allows spatial encoding to occur in 2 or 3 dimensions as opposed to the only 1 covered by the line-by-line read-out method. This technique allows obtaining a full slice (2D EPI) or volume (3D EPI) after each excitation. However, the main downside of EPI is that it induces distortions in the phase encode direction, caused by the loss of signal occurring during read-out. Distortion can be minimised by reducing the acquisition time (e.g. smaller field of view, parallel imaging), or corrected for by acquiring a field map.

3.1.5 Pulse sequences and image contrast

A pulse sequence describes the series of events (RF pulses, gradient activation, data acquisition) required to acquire MR images ([Figure 3.4](#)). Two important components

are the echo time (TE) and the repetition time (TR). The TE is the time between the excitation pulse (first RF pulse) and the middle of the acquisition, and the TR is the time between repeated excitation pulses.




Figure 3.4: Example Spin-Echo pulse sequence and transverse magnetisation timecourse. In this example T2-weighted Spin-Echo pulse sequence, a refocusing RF pulse induces a recovery of magnetisation in the transverse plane. Differences in transverse magnetisation during the acquisition are responsible for the image contrast. The echo time (TE) represent the time between the excitation pulse and the middle of the acquisition.

TE and TR are important parameters for determining the signal difference between tissue types, namely the image contrast. As mentioned in section 3.1.3, tissue composition influences T1 and T2 relaxation times. If the acquisition time occurs at the time of greatest differences in transverse magnetisation between tissue types, then the image is called T2-weighted. Alternatively, T1-weighting is achieved by shortening the TR so that the longitudinal magnetisation does not have time to recover in-between two excitation pulses, which will reveal differences in T1 relaxation between tissue types.

3.2 Diffusion-weighted Imaging

Diffusion-weighted Imaging (DWI) is a non invasive technique that provides an indirect visualisation of the white matter microstructure of the brain. It relies on the measure of water diffusion in different tissue types.

3.2.1 What is diffusion?

In the context of MRI, diffusion refers to the movement of water molecules (Brownian motion), which are contained in high concentration in the brain (Whittall *et al.* 1997). In an unrestricted context, water molecules are equally likely to move in all directions; their diffusion is called *isotropic* (Figure 3.5A). Alternatively, in the presence of restricting boundaries, water diffusion is called *anisotropic*, i.e. the movement of water molecules occurs preferentially along certain directions. Because of their high lipid content (O’Brien & Sampson 1965), myelinated axons are strongly impermeable, and, therefore, strongly restrict intra- and extra-cellular water diffusion to their main direction (Figure 3.5B). This principle is exploited in DWI to infer white matter microstructure properties based on the water diffusion measured empirically in the living brain (Hagmann *et al.* 2006).




Figure 3.5: Water diffusion in brain tissue. Water diffusion (represented by the red 2D tensor) is constrained to the main direction of axons.

3.2.2 Diffusion-weighted MRI sequences

As described in section 3.1.3, a 180° refocusing RF pulse can be used to cancel out the dephasing caused by local magnetic field inhomogeneities, and regain transverse magnetisation (spin-echo sequence). If the refocusing RF pulse occurs at a time t and if protons' resonance frequency remains constant, then spins will be back to their original phase at a time $2t$. Alternatively, if the resonance frequency of protons changes during the post-excitation time (i.e. in the time between the excitation RF pulse and the acquisition), rephasing will only be partial, which will induce a signal loss at the time of acquisition (Figure 3.6). This phenomenon is used in DWI to measure the extent of diffusion occurring between the excitation pulse and the acquisition.




Figure 3.6: Diffusion-weighted spin echo sequence. Transverse magnetisation is represented over the course of a TR in 3 situations. If water molecules are stationary (grey line), rephasing will be complete and there will be no signal loss at the time of the acquisition. Alternatively, moderate (blue line) or rapid (red line) diffusion along the position encode/decode direction will result in an incomplete rephasing, and hence a signal loss proportional to the amount of diffusion.

Figure 3.6 illustrates a typical spin-echo DWI pulse sequence. It includes a posi-

tion encode gradient, active between the excitation RF pulse and the refocusing RF pulse, that modifies the resonance frequency as a function of space (see [Equation 3.1](#)). The role of this gradient is to induce different amounts of dephasing depending on the position of protons in space. Hence, if a proton moves through space during this time, it experiences a unique amount of dephasing dictated by its trajectory. After the refocusing RF pulse occurs, the same gradient, now called position decode gradient, is applied a second time. In a given voxel, if protons have remained stationary (no diffusion) or have not moved along the direction of the position encode/decode gradients, the amount of rephasing will be maximal, i.e. no signal loss occurs. Alternatively, if protons have moved along the direction of the position encode/decode gradients, rephasing will be limited due to experienced changes in resonance frequency, i.e. there is a signal loss proportional to the magnitude of diffusion that occurred between the excitation and acquisition. At every TR of a DWI acquisition, the direction and magnitude of the position encode/decode gradients are varied in order to measure diffusion along as many directions as possible. Typically, DWI sequences use an EPI read-out method in order to acquire a full volume at every repetition.

3.2.3 Image artefacts: distortions and eddy current

Because most DWI sequences have an EPI read-out, they suffer from distortions caused by signal loss occurring throughout the acquisition. Distortions can be limited by increasing the speed at which volumes are acquired (e.g. parallel imaging, partial k-space, smaller field of view). Acquiring two volumes with opposite distortions (so called "blip-up, blip-down"), allows correcting for distortions at the image processing stage using the FSL tool TOPUP (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/topup>) (Andersson *et al.* 2003; Smith *et al.* 2004).

DWI sequences require strong position encode/decode gradients that take time to switch on and off. This can result in the position decode gradient to be fading out during the data acquisition, hence creating some distortions. Contrary to the ones described above, these distortions are direction-specific, i.e. every repetition suffers from distortions along a different direction. This type of eddy current artefact can be accounted for in the preprocessing stage using the FSL tool EDDY (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/eddy>) (Andersson & Sotiropoulos 2016).

3.2.4 Diffusion Tensor Imaging and derived scalars

The result of a DWI acquisition is a 4-dimensional map where the fourth dimension represents the MR signal intensity obtained for any given value and direction of the position encode/decode gradients. One challenge of DWI processing is to reduce the dimensionality of the data using a mathematical representation of diffusion (Hagmann *et al.* 2006). Diffusion tensor fitting is a widely used way of representing diffusion measured empirically in each voxel as a 3D ellipsoid. For any given voxel, acquisitions along all directions are summarised using only 6 values: 3 eigenvectors (ϵ) that define the orientation of the ellipsoid, and 3 eigenvalues (λ) that characterise the extent of diffusion occurring along these directions (Figure 3.7). Anisotropic diffusion is characterised by one eigenvalue being much larger than the other two (e.g. $\lambda_1 > \lambda_2$ & $\lambda_1 > \lambda_3$). Alternatively, if diffusion is isotropic, the tensor is a sphere and all three eigenvalues have similar values (i.e. $\lambda_1 \approx \lambda_2 \approx \lambda_3$).

It is possible to extract several useful scalar measures from diffusion tensor images that can be used as indirect markers of white matter microstructure. For example, Mean Diffusivity (MD) is computed by averaging the three eigenvalues ($MD = (\lambda_1 + \lambda_2 + \lambda_3)/3$) and provides a measure of the average water diffusion, independent of any directionality. Fractional Anisotropy (FA) is another commonly




Figure 3.7: Diffusion tensor. **A.** Diffusion is represented as a tensor, characterised by 3 eigenvalues λ and 3 eigenvectors ϵ . **B.** In the ventricles, water molecules can diffuse freely so tensors look like spheres. In the corpus callosum, diffusion is strongly restricted by the compact white matter fibres so the tensors look like elongated ellipsoids.

derived scalar that ranges between 0 and 1 and represents how directional diffusion is, independent of its magnitude. A FA value of 0 means that diffusion is fully isotropic whereas a FA value of 1 means that diffusion is extremely restricted and only occurs in one direction. These diffusion markers are useful because they reduce the dimensionality of diffusion tensor images (6 dimensions per voxel) to a single 3D map that can more easily be used for statistical analysis.

Myelin and cellular membranes are two major factors restricting water diffusion in neural tissue and they can be affected by many biological phenomena. Therefore, the exact relationship between Diffusion Tensor Imaging (DTI)-derived scalar measures (e.g. FA, MD) and underlying white matter microstructure is still a matter of debate (Le Bihan & Johansen-Berg 2012; Johansen-Berg & Behrens 2013; Assaf *et al.* 2017). For example, FA has been shown to be sensitive to the orientation, degree of myelination, diameter, density, and coherence of fibre bundles (Beaulieu 2009). It increases with increasing packing density or decreasing axon diameter (Takahashi

et al. 2002), and decreases with reductions in myelin (Gulani *et al.* 2001). MD on the other hand was first used as an early marker of brain ischemia (Moseley *et al.* 1990; Warach *et al.* 1992). It decreases at the hyperacute stage due to the presence of a cytotoxic oedema, and increases at the chronic stage because of neuronal loss and gliosis (Fung *et al.* 2011). The variety of factors influencing these derived scalars makes the interpretation of the cause of observed changes delicate. Nevertheless, taken all together, diffusion parameters offer useful, non-invasive, indirect measures of local tissue microstructure (Assaf *et al.* 2017).

3.3 Functional Magnetic Resonance Imaging

Functional MRI (fMRI) is a technique used to detect brain activity while individuals are at rest (resting-state fMRI) or performing a task (task-based fMRI) in the MRI scanner. To date no MR technique has been able to directly measure neuronal firing. Instead, the detection of brain activity relies on the detection of an elevation in blood oxygenation associated with increased energy consumption when neurons are active (Ogawa *et al.* 1992).

3.3.1 Haemodynamic response and blood oxygen-level dependent (BOLD) signal

The generation and propagation of action potentials requires neurons and glial cells to use and produce Adenosine triphosphate (ATP), a molecule whose synthesis involves the oxydation of nicotinamide adenine dinucleotide (NADH) under aerobic conditions. This increased oxygen demand during neuronal activity is followed within a few seconds by a local increase in blood flow and blood volume. This phenomenon, called haemodynamic response, results in an immediate increase of the oxygenated-to-deoxygenated haemoglobin ratio in the blood vessels perfusing the active neural

tissue (Fox & Raichle 1986; Hoge *et al.* 1999; Magistretti *et al.* 1999). It is this increase in oxygenated-to-deoxygenated haemoglobin ratio that is the basis of the Blood Oxygen-level Dependent (BOLD) signal detected with functional Magnetic Resonance Imaging (fMRI).

Haemoglobin is an iron-containing metalloprotein whose magnetic properties vary depending on its binding to oxygen. The loss of oxygen from oxyhaemoglobin (Hb) to deoxyhaemoglobin (dHb) causes the iron to become more paramagnetic, hence inducing more magnetic field inhomogeneity when experiencing an external magnetic field. Under such conditions, the Larmor equation ([Equation 3.1](#)) predicts an increase in dephasing resulting in a shorter $T2^*$ (see section [3.1.3](#)). In summary, following neuronal activity, a local increase in blood oxygenation level induces a reduction in magnetic field inhomogeneities resulting in the magnetisation within a voxel to become more coherent (less dephasing). This phenomenon causes an increase of $T2^*$ that can be measured directly as an increase in MR signal using a Gradient Echo sequence (see section [3.1.4](#)).

3.3.2 fMRI acquisition and artefacts

By far the most common way of acquiring fMRI data is to use a $T2^*$ -weighted gradient echo sequence with an EPI read-out (see sections [3.1.3](#) and [3.1.4](#)). Contrary to spin echo sequences, gradient echo sequences are sensitive to $T2^*$ relaxation because they do not include a refocusing RF pulse. The TE should be chosen within a time window that maximises the signal difference (contrast) between an active (high Hb/dHB ratio, i.e. low magnetic field inhomogeneity, i.e. high signal) and a passive (low Hb/dHB ratio, i.e. high magnetic field inhomogeneity, i.e. low signal) condition ([Figure 3.8](#)). For grey matter, this optimal TE range falls around 45 ms at 3 Tesla and 25 ms at 7 Tesla.




Figure 3.8: BOLD signal timecourse and optimal TE. This graph illustrates the timecourse of transverse magnetisation under an active (red) and resting (dotted blue) conditions. In green is represented the difference between the two conditions (scaled). The time of greatest difference between the active and rest MR signal defines an optimal range for the acquiring data (TE).

fMRI images typically suffer from dropouts and distortions. Dropouts result from B_0 inhomogeneities caused by air-tissue boundaries. Decreasing the TE minimises dropouts at the cost of reducing the BOLD sensitivity (i.e. less contrast between the active and rest conditions, [Figure 3.8](#)). Alternatively, acquisition parameters such as slice orientation and slice thickness can also be modulated in order to minimise dropouts. Image distortion is an EPI artefact occurring in the phase encode direction (see [section 3.2.3](#)) that can be reduced by increasing the speed of data acquisition (e.g. parallel imaging, partial k-space, smaller field of view). Image distortions can also be partially corrected at the pre-processing stage using a field map, an image of the magnetic field inhomogeneities.

3.3.3 Statistical analysis of task fMRI data

The aim of task-based fMRI experiments is to identify brain regions that are active while individuals perform a task, and characterise their contribution to a behaviour

of interest. The most common way to analyse task fMRI data is to fit a general linear model (GLM) to the signal of each voxel in order to identify brain regions whose BOLD signal is linearly associated with the task (Smith *et al.* 2004). First, a predicted BOLD response is calculated based on the convolution of a haemodynamic response function (for example Buxton *et al.* 2004; Shan *et al.* 2014) and predicted neural activity. The predicted neural activity is derived from the task the participant is performing in the scanner during the fMRI acquisition, and requires the experimenter to build a model of how neurons should respond to certain stimuli. The predicted BOLD response can then be fitted to the measured BOLD response of each voxel within a GLM framework in order to compute a t-statistic that represents the ratio between fitted amplitude and residual noise. Typically, voxelwise t-statistic maps are then thresholded and corrected for multiple comparisons. These maps can then be analysed at a higher level to evaluate differences between conditions or groups for example.

Because of the temporal properties of the haemodynamic response function, two stimuli should be spaced by at least 2 seconds in order to be distinguishable. Hence, there are two main classes of experimental designs that differ in terms of the timing of the task: block design and event-related design. Block design experiments include active blocks during which a task is performed continuously (at a high trial frequency) and blocks of rest and/or blocks of a control task. Typically, individual trials cannot be distinguished and instead the predicted neural activity has the shape of a box lasting the duration of an active block (for example, see second row of [Figure 3.9](#)). Typically, brain regions associated with a particular psychological process (e.g. psychological process X) are isolated by contrasting experimental blocks in which the process is involved (e.g. task 1 involves psychological processes X Y Z) to control blocks that theoretically involve the same processes except the one of interest (e.g. task 2 requires psychological processes Y Z). This technique relies on the assumption




Figure 3.9: General Linear Model (GLM) The haemodynamic response function (in green) is convolved with a predicted neural response (in black) in order to obtain a predicted BOLD response (in blue) that is fitted to the BOLD timeseries of each voxel in the image (in red). Two example GLM fits are shown for illustration purposes. The top one yields a high t-statistic because the residual noise is low relative to the fitted magnitude. The bottom one yields a low t-statistic because the residual noise is high in relation to the fitted magnitude.

that processes can be neurally and psychologically *added* ('pure insertion' principle) and implies no interaction among them (Amaro & Barker 2006). Alternatively, factorial, parametric and conjunction block designs offer a more precise characterisation of the interaction between psychological processes. Event-related designs require individual stimuli to be spaced in time in order for the predicted BOLD response to capture each of them separately. When coupled to an algorithmic account of behaviour, this type of experimental design enables a more precise characterisation of the computations implemented within the active neural tissue (for example Daw *et al.* 2006; Hampton *et al.* 2006; Kim *et al.* 2015).

3.3.4 Statistical analysis of resting state fMRI data

The aim of resting state fMRI studies is to infer functional coupling between brain regions from spontaneous BOLD fluctuations measured *in vivo*. Typically, BOLD fMRI is acquired under conditions of limited sensory and cognitive stimulation where volunteers are instructed to lie still and either fixate on a cross hair projected on a screen, or lie awake with their eyes closed for 5 to 10 minutes. During this time, BOLD signal is thought to reflect spontaneous haemodynamic fluctuations in the brain but it is also contaminated by artefacts such as scanner instability or non-neuronal physiological noise (e.g. breathing, cardiac pulse, head motion, etc.) (Soares *et al.* 2016). Several methods have been developed to tease apart physiological signal from noise. Note that these methods can also be applied to task fMRI data. The first method requires independent measures of potential sources of noise in order to regress out their influence on the total BOLD signal (Birn *et al.* 2006; Murphy *et al.* 2013). The second method uses data-driven approaches to identify spatio-temporal signatures of noise signal. For example, regressing out the signal that is common across all voxels (Macey *et al.* 2004) or that originates from regions strongly affected by physiological artefacts (Zarahn *et al.* 1997) have been used for this purpose. More recently, manual or automatic detection and removal of noise components following independent component decomposition has become the gold-standard for BOLD data denoising (Griffanti *et al.* 2014; Pruim *et al.* 2015). This method will be described in more detail in the following paragraphs.

Analysis of resting BOLD data relies on the identification of networks of brain regions whose BOLD signal is temporally correlated. This can be done in a hypothesis-driven way where BOLD signal timecourses are extracted from one (seed-based analysis) or more (cluster-based analysis) region(s) of interest and regressed against the rest of the brain or against each other. Although simple and relatively easy to in-




Figure 3.10: Independent Component Analysis (ICA). Schematic representation of the data structure and decomposition performed by ICA on BOLD fMRI data. Figure adapted from Beckmann (2012)

terpret, this approach relies heavily on a priori definition of regions of interest and has little exploratory power. Furthermore, since only a subset of the BOLD data is considered, it does not allow for multiple systems to be studied simultaneously. Independent component analysis (ICA) offers a solution to these limitations by summarising the entire BOLD dataset into a set of statistically independent voxelwise spatio-temporal components (McKeown *et al.* 1997; Beckmann *et al.* 2005; Smith *et al.* 2009; Beckmann 2012). In brief, this dimensionality reduction technique can be seen as a GLM in which no design matrix (predicted neural activity) is specified. The BOLD data is first reorganised into a 2D 'time $\times$ space' matrix by re-arranging the intensity of all voxels for a given time point into a single row (Figure 3.10). It is then decomposed into the product of two matrices as follows: 1) a 'time $\times$ component' matrix containing a timecourse of the underlying signal for each component; 2) a 'component $\times$ space' matrix containing the spatial map of each component. Note that this technique is a way of finding temporal basis vectors so that the associated spatial maps are sparse and statistically independent. To come back to the GLM analogy, the 'time $\times$ component' matrix takes the role of the GLM design matrix while the 'component $\times$ space' represents maps of the effect size (Beckmann 2012). The main difference here is that both are estimated simultaneously from the data, using information theoretic principles. ICA analysis can be run in FSL using the

Multivariate Exploratory Linear Optimised Decomposition into Independent Components (MELODIC) tool (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/MELODIC>) (Beckmann & Smith 2004).




Figure 3.11: Resting State Networks (RSNs). Image of the 10 most representative RSNs (on the left of each panel), matched to their corresponding functional BrainMap (BM) networks (on the right of each panel). Figure adapted from Smith *et al.* (2009)

Typically, some of the components extracted with ICA reflect noise (Griffanti *et al.* 2014) and can be removed from the data. Others, however, resemble discrete cortical functional networks (Beckmann *et al.* 2005; Smith *et al.* 2009) (Figure 3.11). These components, called resting state networks (RSNs), are thought to reflect *functionally* connected brain regions (Beckmann *et al.* 2005) and are only partially constrained by structure (Honey *et al.* 2009; Hermundstad *et al.* 2013). Figure 3.11 illustrates the most commonly reported RSNs (on the left) matched to their corresponding functional networks from the BrainMap database (on the right) (Smith *et al.* 2009), and shows a large degree of overlap between the two. Recently, this coupling between brain connectivity and function has even been captured at the individual-subject level

(Tavor *et al.* 2016). Of particular interest to this thesis are the sensorimotor network, including the pre- and post-central gyri and the SMA (RSN 6₂₀ in [Figure 3.11](#)), and the frontoparietal networks, including the bilateral intraparietal sulcus and right or left middle and superior frontal gyri (RSNs 9₂₀ & 10₂₀ in [Figure 3.11](#)) (Beckmann *et al.* 2005; Smith *et al.* 2009) (see [chapter 5](#) & [7](#)).

3.4 Magnetic Resonance Spectroscopy

Magnetic Resonance Spectroscopy (MRS) is a non invasive technique that allows the quantification of endogenous metabolites in a localised region of interest. The output of the acquisition is a spectrum (signal intensity as a function of resonance frequency) that contains signal from many different metabolites, each of them identifiable by a unique and highly reproducible frequency distribution. Knowing the specific signatures of these metabolites allows the measurement of their contribution to the final spectrum and the extrapolation of their concentration in the region of interest.

3.4.1 Principles of *in vivo* MR spectroscopy

We described in section [3.1](#) the magnetic resonance effect on the positively charged spinning nucleus of hydrogen. We saw that the Larmor equation ([Equation 3.1](#)) describes a linear relationship between nuclei's resonance frequency ω_0 and the strength of the magnetic field B_0 they are experiencing. However, the latter is also influenced by the local chemical environment when hydrogen atoms form a chemical bound. This effect, called chemical shielding, is responsible for a shift of the resonance frequency and requires the Larmor equation to be re-written as follows:

$$\omega_0 = \gamma \times B_0(1 - \sigma) \quad (3.2)$$

where σ is a chemical shielding constant (bounded between 0 and 1) that depends on the shape of the electron cloud surrounding the nucleus. Because of differences in chemical and geometrical composition, every molecule exerts a unique chemical shielding on its hydrogen nuclei, resulting in a unique resonance frequency (or chemical shift) that can be used for MR spectroscopy fingerprinting of chemical compounds. In MRS, resonance frequency is typically expressed as a ratio of the resonance frequency of an isolated proton, in parts per million (ppm), in order to calculate a metric that is independent of the B_0 strength and therefore more readily comparable across scanners and acquisitions.

In summary, when interpreting a MRS spectrum, the position of the peaks on the frequency axis determines the amount of shielding experienced by the hydrogen nuclei, and can therefore be used as a signature of chemical compounds. The amplitude of these peaks is proportional to the number of hydrogen nuclei experiencing a certain shielding, and provides information about the metabolite concentration (see section [3.4.3](#)).

3.4.2 Acquisition considerations: water suppression, outer volume suppression and shimming

As described in section [3.2](#), hydrogen nuclei contained in water constitute the major source of MR signal in the brain. If unsuppressed, water would create a large peak (at 4.65 ppm) overlapping almost all relevant metabolite peaks and making the quantification of these metabolites impossible. Water suppression refers to the process of minimising the water signal by applying a frequency selective 90° RF pulse prior to the excitation. The goal of this RF pulse is to selectively flip the magnetisation of water protons to the transverse plane in order to make them insensitive to the upcoming excitation pulse. A similar procedure, now called 'outer volume suppression',

can be applied to the lipid signal (1.3 ppm) originating from the dura, cranial bone and skin. This is of particular importance for primary motor cortex MRS because of the proximity of the region of interest to the surrounding non-brain tissue.

Shimming is the process of optimising the homogeneity of the magnetic field B_0 over the region of interest. It is critical for MRS because it directly narrows the spectral linewidth, making separation and quantification of metabolites more reliable (Juchem & de Graaf 2017).

3.4.3 Standard spectrum analysis and metabolite quantification

The experimentally measured spectrum combines information from all metabolites present in the volume of interest. Metabolite quantification relies on the decomposition of this spectrum into its constituent metabolites and typically involves fitting a scaled linear combination of the potential individual metabolite contributions (Provencher 2001) (Figure 3.12). Thus, the analysis of spectra requires prior knowledge of the individual spectral fingerprint of potentially detected metabolites. This information, called basis-set, is typically obtained experimentally by measuring metabolites of known concentration in solution using the same acquisition parameters as the *in vivo* experiment (i.e. same pulse sequence, same B_0 field strength).

Metabolite quantification is typically expressed in reference to a metabolite of well-known concentration. There are two main approaches to referencing for MR spectroscopy: external and internal methods (Stagg & Rothman 2013). The external referencing method involves, for each MRS scan, the acquisition of a second MRS volume of interest centred on a small vial placed near the individual's head and containing a metabolite in known concentration. The advantage of this method is




Figure 3.12: Components of a spectrum. A typical model spectrum is shown on the top row (dark blue). The individual fits for the main metabolites are shown in the following rows. Metabolites with low concentrations (e.g. GABA) are reflected by the relatively small amplitude of their peaks.

that the concentration of the reference metabolite is constant and precisely known. However, because the reference signal comes from a different region than the *in vivo* volume of interest, it is affected by different magnetic field inhomogeneities, which may confound its relationship to the signal of interest. Internal referencing offers an alternative to this issue. It involves choosing a metabolite, whose concentration is stable over time and across individuals, to serve as reference compound. The most commonly used metabolite for this purpose is total Creatine (Creatine + phosphocreatine) whose concentration is relatively stable around 8 mM in the healthy brain (Provencher 2012). If there are reasons to believe Creatine concentration is not stable (e.g. due to certain pathological processes), its use for internal referencing is inappropriate. Compared to the external referencing method, internal referencing has the advantage of involving a signal originating from the same volume of interest as

the spectroscopy measurement, which comes at the cost of not knowing the precise concentration of the reference compound.

3.4.4 Quality filtering

3.4.4.1 Comment on the use of the relative Cramér-Rao Lower Bounds

Because of the noisy nature of MRS data, researchers have tried to design a method to systematically identify bad quality data in an objective way that does not depend on visual inspection by an expert. Over the years, the method of Cramér-Rao Lower Bounds (CRLB) has become the gold standard for determining concentration estimate uncertainty (Cavassila *et al.* 2001). The CRLB represents the lowest possible standard deviation of all unbiased concentration estimates obtained from fitting the model to the data. The CRLB is broadly used in its *relative* form — labelled *%CRLB* — as a percentage of the estimated metabolite concentration. Its calculation for all estimated metabolites is part of the standard analysis pipeline of LCModel (Provencher 1993; Provencher 2001). Many authors have used the *%CRLB* as a way to perform quality filtering by rejecting metabolite concentration estimates for which the *%CRLB* falls above a certain arbitrary threshold judged as unacceptable level of uncertainty in the estimate (usually between 20% and 50%) (Rowland *et al.* 2012; Stone *et al.* 2012; Kim *et al.* 2014; Bachtiar *et al.* 2015; Barron *et al.* 2016).

Recently, some authors have warned against the usage of *%CRLB* for quality filtering because it could lead to wrong or missed statistical findings (Tisell *et al.* 2013; Kreis 2016). For equivalent metabolite concentrations, large levels of noise caused by bad quality acquisition (e.g. too small voxel size, not enough averages, bad shimming) or bad quality spectrum fitting (e.g. inappropriate basis files) would result in an increase in *%CRLB*, in a way that truly reflects estimation uncertainty. Thus, in this scenario, it would be valid to mistrust the data based on a high *%CRLB*. However, because of the *relative* nature of the *%CRLB*, this metric also strongly de-




Figure 3.13: Relationship between concentration estimate and relative Cramér-Rao lower bounds. Because of the nature of the relative CRLB, lower concentration estimates tend to lead to greater %CRLB values. The red line shows the 'constant noise model' fitted to GABA (A) and Glutamate (B) concentration estimates for data points that have a %CRLB lower than 100%. The color code expresses, for each data point, the distance from the 'constant noise model' in terms of Pearson residual. A Pearson residual greater than 2 (i.e. outside the dotted bound), is considered as an outlier.

depends on its denominator, i.e. the estimated metabolite concentration. Thus, for equivalent levels of noise, a true decrease in the metabolite concentration would also be associated with an increase in %CRLB. In this scenario, however, it is no longer valid to reject such datasets. To illustrate this point, Figure 3.13 shows the relationship between concentration estimate (x-axis) and %CRLB (y-axis) for the 96 MRS datasets acquired on a 7T Siemens MAGNETOM scanner at the Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB) in the study reported in chapter 5. Visual inspection of the scatter plots for GABA (Figure 3.13A) shows that the rejection of datasets with a %CRLB greater than 50% (data points above the black line) would induce a selection bias towards high concentration estimates. When measuring a change in GABA concentration following an intervention (e.g. anodal transcranial direct current stimulation, a-tDCS), this selection bias may artificially inflate the chances of detecting a reduction post-intervention, simply by virtue of regression towards the mean (Blance *et al.* 2005; Tu & Gilthorpe 2007). Additionally, it

might be the case that individuals with higher baseline GABA levels are more likely to respond to a-tDCS than those with low basal GABA levels.

3.4.4.2 Is there a better alternative?

We propose that this methodological caveat can be avoided by taking into consideration the concentration estimate when rejecting datasets with high $\%CRLB$. Datasets might have a high $\%CRLB$ because of a low concentration estimate rather than an excessive level of noise. We suggest that such datasets should not be excluded. Alternatively, datasets might have low $\%CRLB$ merely because the concentration estimate is high. However, it might be the case the the level of $\%CRLB$ is excessively high, given the metabolite concentration. Such datasets should be excluded.

We propose the following method as an alternative to standard $\%CRLB$ -based quality filtering. First, the following model is fitted to the 'concentration estimate $\times$ $\%CRLB$ ' relationship:

$$Expected\ \%CRLB_i = \frac{N_i}{C_i} \quad (3.3)$$

where N_i represents a group noise constant and C_i the concentration estimates for a metabolite i . If this simple model can explain most of the variance in the observed relationship between concentration estimates and $\%CRLB$ (Figure 3.13), it means that the level of noise is relatively constant across all measurements. However, any deviation from this model reflects an *unusual* level of noise compared to the other measurements. For each measurement, deviation from the model can be expressed as the Pearson residual as follows:

$$e_i = \frac{r}{\sqrt{MSE}} \quad (3.4)$$

where r_i is the raw residual (i.e. difference between the $\%CRLB$ and *expected* $\%CRLB$ for a certain measurement) and MSE is the mean squared error (i.e. mean deviation of all measurements from the model). The greater the Pearson residual for a given measurement, the noisier it is in regard to the rest of the data, irrespective of the concentration estimate. [Figure 3.13](#) illustrates measurement rejection based on a threshold of 2 for this metric. Note that this method did not reject the lower tail of the distribution entirely and therefore does not induce a selection bias towards high concentration estimates. We used this method in [chapter 5](#) and [6](#).

Chapter 4

Inferring computational states driving adaptation dynamics

In this experimental chapter, we modify an influential two-state model of sensorimotor adaptation (Smith et al. 2006a), which significantly improves its ability to describe and predict motor behaviour during and after prism adaptation, across all experimental groups in this thesis (dataset 1: $n = 25$ young adults; dataset 2: $n = 24$ young adults; dataset 3: $n = 23$ elderly adults; dataset 4: $n = 15$ neglect patients). The model posits two internal adaptive systems that learn and forget at different rates, and that combine to explain motor behaviour during adaptation. By adding a weighting factor, to allow for a variable contribution of each learning system to the expression of prism after-effects, our modification improved the quality of the model fit, and made it possible to determine the relative weight individuals assign to the learning from each system, which in turn determines the decay rate of the prism after-effect. The theoretical model described here is of central importance to much of the work presented in this thesis.

4.1 Introduction

The exposure of individuals to a novel sensorimotor environment – one that requires a new mapping between sensory information and motor actions – initially induces large directional errors in reaching movements. However, if sufficient practice is allowed, individuals are able to compensate for such perturbation (i.e. correct their reaching movements) so as to gradually regain baseline performance (Shadmehr *et al.* 2010). For example, when wearing prism glasses that displace the visual field towards the right, individuals initially make large rightward errors on visually-guided reaching movements. Within a few practice trials (typically 10-20), healthy individuals can adapt to the visual displacement and re-gain baseline accuracy (i.e. reach visual targets accurately). This phenomenon is called 'prism adaptation' (PA) (von Helmholtz 1867). If the prism glasses are then removed, participants miss the visual target in the opposite direction, i.e. towards the left (Redding & Wallace 2002; Redding *et al.* 2005). This phenomenon is called 'after-effect' (AE), and reflects the temporary persistence of some of the compensatory mechanisms engaged during the adaptation (Harris 1963).

PA is of particular interest for rehabilitation of post-stroke spatial neglect (Rossetti *et al.* 1998; Barrett *et al.* 2012), a syndrome characterised by the failure to report, respond or orient to novel or meaningful stimuli presented to the contralesional (left) side (Heilman & Van Den Abell 1980; Heilman *et al.* 2002). The acquired leftward bias following PA (i.e. prism AE) appears to generalise abnormally in neglect patients, leading to improvements in untrained cognitive domains (e.g. visual search, reading, wheelchair driving, mental representation) (for review, see Jacquin-Courtois *et al.* 2013). The induction of a sensorimotor AE (i.e. leftward shift in pointing behaviour) is believed to be an essential prerequisite to the therapeutic improvements of neglect following PA (Serino *et al.* 2006; Newport & Schenk 2012;

Jacquin-Courtois *et al.* 2013; Rode *et al.* 2015). Maximising persistence of the AE indeed appears to be a way of inducing more persistent therapeutic effects (O’Shea *et al.* 2017). Hence, it becomes crucial for the field to better understand the AE memory trace, the conditions under which it persists, its underlying neural mechanisms, and how to influence it in favour of long-term retention. This is the central goal of this thesis. The present chapter aims to define a principled, quantitative, generative framework within which the AE memory trace can be studied.

Computational theories of motor control posit that when a perturbation (such as an optical shift) causes reaching errors, the brain forms an internal model of the perturbation, used to counteract the disturbance and restore accuracy (i.e. adapt) (Shadmehr *et al.* 2010; Wolpert *et al.* 2011). This type of learning is often termed model-based learning (Shadmehr *et al.* 2010). In the context of sensorimotor adaptation, it was demonstrated that the internal models the brain forms are intrinsically stamped with a decay rate, which explains AE persistence (Körding *et al.* 2007; Joiner & Smith 2008; Inoue *et al.* 2015). Depending on what is believed to cause performance error, the solution of the adaptation will either be retained or forgotten, leading to stable or labile AEs. These hidden internal states driving behaviour can be inferred directly from individuals’ behavioural data using computational tools such as state-space or Bayesian models (Kitazawa *et al.* 1995; Körding & Wolpert 2004a; Smith *et al.* 2006a; Körding *et al.* 2007; Berniker & Körding 2008; Joiner & Smith 2008; Haith *et al.* 2009; Lee & Schweighofer 2009; Yamamoto & Ando 2012; Inoue *et al.* 2015; McDougle *et al.* 2015). One key difference between these two classes of models is that the former extracts hidden internal states from the raw behavioural data whereas the latter infers them from the noise structure of the experiment, under the assumption of optimal behaviour. Bayesian models therefore require the reliability of different sources of information (e.g. vision, proprioception) to be

experimentally manipulated (for examples of such experimental manipulations, see Körding & Wolpert 2004a; Wei & Körding 2010). The set of experiments reported in this thesis was designed to measure the evolving dynamics of the AE during and after adaptation to prisms and did not involve manipulations of the reliability of different sources of information. Hence the state-space modelling framework was used.

Smith and colleagues’ influential two-state model posits that adaptation arises from the interaction of two systems – one fast and one slow – that learn and forget over different timescales (Smith *et al.* 2006a). At any given time point, the sum of their states is thought to provide an internal model of the perturbation that is used to correct the motor output. According to this view, both systems contribute equally to motor behaviour, with the fast system primarily driving error correction during the early phase of prism exposure, and the slow system during the later phase. Here, we hypothesised that individuals might differ in the weight they assign to each system when expressing prism AEs, explaining inter-individual differences in AE persistence. Assigning more weight towards the slow system might predict a more stable AE memory trace. Based on this hypothesis, we built a variant of the two-state model allowing a variable contribution of the fast and slow systems to the prism AE. We asked the following questions:

1. Compared to the original model, does this variant yield a closer fit to PA data, allowing a more accurate estimate of the hidden internal states underlying the AE memory trace?
2. In the immediate period following PA, what is the predictive value of this variant compared to the original model?
3. Does this variant predict the AE memory trace 24 hours after PA?

We demonstrate that this variant outperforms the original model across all experi-

ments reported in this thesis and is able to predict long-term retention of the AE. We therefore propose to use the hidden internal states inferred with this principled, quantitative, generative framework to understand the AE memory trace in the rest of this thesis.

4.2 Material and Methods

4.2.1 Prism adaptation protocol

The PA paradigm described in this section was used in experiments reported in chapters 5, 6 and 8, so it is important to clarify in detail the rationale and procedures. Most PA protocols used in the literature include a long prism exposure block followed by a measure of the after-effect (Martin *et al.* 1996; Redding & Wallace 1996; Fernández-Ruiz *et al.* 2000; Schintu *et al.* 2014; Hanajima *et al.* 2015) (Figure 4.1A). This type of protocol is not well suited for computational modelling because error-dependent learning asymptotes rapidly after the initial error correction phase, and pointing behaviour remains at the same accuracy level for the rest of the exposure. However, a lot of what determines AE persistence happens during this phase of stable behaviour, and is therefore invisible in this type of behavioural data. By contrast, the PA protocol used in this thesis was designed to assay the AE and its tendency to decay or persist at various time points throughout PA (Inoue *et al.* 2015; O’Shea *et al.* 2017). To do so, interleaved measures of AE are taken throughout prism exposure (Figure 4.1B). Data collected with this protocol therefore provide a direct measure of the evolving dynamics of the AE as adaptation develops. Crucially, this enables a better modelling of the hidden internal states underlying the AE memory trace, a central goal of this thesis.




Figure 4.1: Prism adaptation protocol. **A.** Schematic representation of the most common prism adaptation (PA) protocol. Typically, accuracy on closed-loop (CLP, black blocks) and open-loop (OLP, red blocks) is first measured at baseline. Then follows a prism exposure phase during which participants repeatedly point at visual targets while wearing prisms. Visual feedback of reaching error drives error-corrective processes. The after-effect (AE) is finally measured at the end of the experiment. **B.** Schematic representation of the PA protocol used in this thesis. The key difference is the adaptation phase that includes blocks of CLP, during which participants adapt to prisms, interleaved with blocks of OLP, used to measure the AE. **C.** Example behavioural performance in a single participant. The x-axis represents trial number and the y-axis pointing accuracy (rightward errors are positive). Black wedges indicate blocks throughout which prisms were worn. During Adaptation (prism exposure, E1-6), participants see the outcome of each trial, so they can correct their errors from trial-to-trial. The AE is measured in interleaved blocks without visual feedback (AE1-6, shaded light grey). This participant shows the stereotypical behavioural pattern observed in all experiments. CLP was right shifted early during the Adaptation but was gradually corrected. As a result, OLP became more left shifted (AE).

The PA protocol was as follows. Participants sat with their chin positioned on a chinrest and were instructed to perform, with their right hand, centre-out reaching movements to one of three targets situated in front of them. The distance between

their eyes and the central target was approximately 57 cm. The two lateral targets were situated 10 cm to the left and right of the central target. Individuals performed two types of trial: closed-loop pointing (CLP) and open-loop pointing (OLP).

Closed-loop pointing. On CLP trials, participants made speeded reaching movements to point with their right index finger to the left and right targets, in pseudo-random order. Both accuracy and speed were emphasised. Similar to previous experiments (O’Shea *et al.* 2014; Inoue *et al.* 2015; O’Shea *et al.* 2017), visual feedback was limited to the last 2/3 of the reaching movement in order to limit strategic adjustments (Redding & Wallace 1996; Redding & Wallace 2001) and in-flight error correction. [Figure 4.1C](#) shows pointing accuracy over time for a participant undergoing the PA protocol. It illustrates the stereotypical pattern of performance observed in all individuals across all experiments. Closed-loop pointing (in shaded red) was accurate at baseline and became right-shifted during the initial prism exposure period due to the visual displacement ([Figure 4.1C](#): E1). Throughout prism exposure (E1-7), participants saw the outcome of each trial, allowing error-dependent learning to proceed. From trial-to-trial, participants gradually corrected their errors until they regained near-baseline accuracy levels ([Figure 4.1C](#): compare E1 to E6).

Open-loop pointing. On OLP trials, participants pointed at a comfortable speed to the central target. The target was only visible at the reach onset. Visual feedback of the reaching movement, terminal error and return movement were occluded by an opaque shutter, hence requiring participants to rely on proprioception to guide their movement. Accuracy was emphasised over speed. The rationale for the occlusion of visual feedback is to enable the AE to be measured without participants receiving that feedback and subsequently de-adapting (i.e. learn from performance error in the opposite direction). Compared to baseline, pointing became left-shifted in

the OLP blocks interleaved between prism exposure CLP blocks (Figure 4.1C: AE1-6). Note that in this individual, the AE is more labile during the early phase of the adaptation (Figure 4.1C: the first trials of AE1 are more left shifted than the last ones, illustrating intra-block decay of the AE) compared to the later phase (Figure 4.1C: pointing error is stable over time in AE6, illustrating little decay of the AE), reflecting consolidation of the AE memory trace throughout prism exposure. Although CLP behaviour is quite stereotypical across participants, the dynamics of AE development show more inter-individual variability. Explaining this inter-individual variability is one of the goals of the computational work described in this chapter.

In all experiments, participants first performed closed-loop and open-loop pointing to measure their baseline accuracy on these two tasks. Next, the adaptation phase consisted of 6 blocks of prism exposure (Figure 4.1C; E1-6: 10 CLP trials in block 1-2, 20 CLP trials in blocks 3-6; equal number of trials to each of two target locations) alternating with 6 blocks of AE measurement (Figure 4.1C; AE1-6: 15 OLP trials per block).

A final important point about this PA protocol is that error-dependent learning and AEs are not quantified under identical conditions. Instead, two different tasks are used. Error-dependent learning is quantified, under prism exposure with visual feedback, as endpoint accuracy on trials of speeded centre-out reaches (CLP) to lateral targets. AE is measured, after prism removal, without visual feedback, as endpoint accuracy under non-speeded conditions (OLP) to an untrained central target. By quantifying that portion of learning that transfers across a contextual change (of limb dynamics, speed, feedback, and spatial location), this AE measure reflects generalisation (Kitazawa *et al.* 1997), a factor thought to be an essential precursor of the cognitive transfer gains from PA observed in neurological patients with visual neglect

(Jacquin-Courtois *et al.* 2013). The relevance of this point for the computational model proposed in this chapter is discussed in the next section.

4.2.2 Modelling framework

4.2.2.1 Existing state-space models of adaptation

Various computational models have been proposed to account for the trial-to-trial temporal dynamics of motor performance during sensorimotor adaptation. One of the earliest attempts is a linear time invariant single-state model with full trial-to-trial retention (Kitazawa *et al.* 1995). This model posits that the endpoint error on the n -th trial $e(n)$ depends on the error on the first pointing with prisms on $e(0)$ and some constant rate B at which error decreases. The model is defined as follows:

$$e(n) = e(0) - B \sum_{i=1}^{n-1} e(i) \quad (4.1)$$

$$0 < B < 1$$

where the pointing error on the n -th trial is linearly related to the summation of errors from the first to the $(n-1)$ -th trial. This model accounts for the exponential nature of error correction on closed-loop pointing during prism exposure but fails to explain AE behaviour on open-loop pointing. It is unclear from equation 4.1, how the model should behave when the disturbance is removed and visual feedback is deprived.

More recently, Smith *et al.* (2006a) proposed a unified model able to explain motor behaviour during exposure to a perturbation and subsequent measures of after-effects. This model posits that trial-by-trial adaptation arises from the interaction of multiple adaptive systems that can be approximated by two systems: a 'fast' one that learns and forgets quickly, which is mainly responsible for the rapid initial error correction, and a 'slow' one that learns and forgets slowly, allowing the long-term retention of AE

(Smith *et al.* 2006a; Joiner & Smith 2008). When applied to CLP data, the original two-state model is formulated as follows:

$$\begin{aligned}
e(n) &= bias_{CLP} + d - X(n) \\
X(n) &= x_f(n) + x_s(n) \\
x_f(n) &= A_f x_f(n-1) + B_f e(n-1) \\
x_s(n) &= A_s x_s(n-1) + B_s e(n-1) \\
0 < A_f < A_s < 1; \quad 0 < B_s < B_f < 1
\end{aligned} \tag{4.2}$$

where the endpoint error on the n -th trial $e(n)$ is modelled as the sum of a bias term $bias_{CLP}$ and a prism effect d from which is removed the sum of the states of the two systems $X(n)$. The bias term is a fixed parameter capturing the natural deviation from target participants show at baseline (i.e. a constant 'irreducible' error that is not attributable to PA). The prism effect d is a fixed parameter that determines the error magnitude obtained on the first prism exposure CLP trial of the Adaptation. Owing to online corrections during reach trajectory with visual feedback (CLP), d is always lower than the true visual displacement, and is typically about 60-80% (Rock *et al.* 1966; Redding & Wallace 2000; Redding *et al.* 2005). The total amount of adaptation $X(n)$ corresponds to the sum of the states of the fast system $x_f(n)$ and slow system $x_s(n)$. These systems produce a trial-by-trial estimate of d based on the visual error feedback. The learning rate B dictates the proportion of error added to the state of a system from one trial to the other. A greater B means faster learning, i.e. larger proportion of error on the $(n-1)$ -th trial corrected for on trial n . The retention rate A determines the speed of the memory decay, i.e. the proportion of state that is being carried over from one trial to the other. A high A for a system means little inter-trial memory decay.

On OLP trials, where the visual disturbance is removed and no visual feedback is available, the original two-state model (Smith *et al.* 2006a) becomes:

$$\begin{aligned}
 e(n) &= bias_{OLP} - X(n) \\
 X(n) &= x_f(n) + x_s(n) \\
 x_f(n) &= A_f x_f(n-1) \\
 x_s(n) &= A_s x_s(n-1) \\
 0 < A_f < A_s < 1
 \end{aligned} \tag{4.3}$$

This equation differs from equation 4.2 in three ways. First, a different bias term is used, reflecting the fact that, at baseline, individuals show different biases between the two tasks (for example, compare the baseline deviation on the CLP and OLP blocks in Figure 4.1C). Second, the prism effect d is removed (effectively set to zero) because individuals are no longer wearing prisms when the AE is measured. Third, the learning rates of the fast and slow systems are set to zero (i.e. $B_f = 0$, $B_s = 0$) to reflect the absence of visual feedback of the reach endpoint error (the signal driving learning of the two systems) during OLP trials. Any within-block change in pointing accuracy during AE measurements is therefore explained by the forgetting of the two systems in this model. If visual feedback is provided during AE measurements (i.e. CLP, e.g. de-adaptation phase of experiment 1), the learning of both systems is switched on and the equation is the same as in 4.2, apart from $d = 0$. In this case, within-block change in pointing accuracy reflects both the forgetting of the adapted state and the active learning from performance error (i.e. de-adaptation).

It is important to note that because this model is linear time-invariant, the learning and retention rates have the same value in all experimental phases (i.e. in equations 4.2 & 4.3) (Smith *et al.* 2006a; Joiner & Smith 2008; Inoue *et al.* 2015). The

underlying assumption is that the fast and slow systems, i.e. the hidden internal computational states driving behaviour, operate at certain dynamics that are un-affected by the passage of time or changes in task constraints.

In summary, when applied to PA data, the original two-state model (Smith *et al.* 2006a) is composed of three fixed parameters (d , $bias_{CLP}$ and $bias_{OLP}$), computed from the raw data, and four free parameters (A_f , A_s , B_f and B_s) estimated via optimisation. Throughout this thesis, we used the same d for all participants of an experiment. Unless stated otherwise, it was calculated as the group mean pointing error on the first CLP trial of the first block of prism exposure, normalised by the mean deviation on the baseline CLP block. The other two fixed parameters ($bias_{CLP}$ & $bias_{OLP}$) were calculated on an individual basis as the block mean pointing error on the baseline CLP and OLP block, respectively.

4.2.2.2 Weighted two-state model

To date, state-space models have been used predominantly in the literature on visuo-motor rotation (Lee & Schweighofer 2009; Tanaka *et al.* 2009; Tanaka *et al.* 2012; Kim *et al.* 2015; McDougle *et al.* 2015) and force-field adaptation experiments (Smith *et al.* 2006a; Joiner & Smith 2008), two paradigms in which the AEs are typically measured using the same experimental constraints as during exposure to the perturbation (i.e. identical effectors, movement dynamics, target location in space). Because of these identical conditions, both error-dependent learning and AE measurements are assumed to reflect the direct sum of the states of the two systems (see equations 4.2 and 4.3) (Smith *et al.* 2006a). However, if the error-dependent learning and AE measurements occur under different experimental conditions (CLP vs. OLP: different target location, different movement speed; see section 4.2.1), the assumption of equal contribution of the two systems to the AE may no longer hold. For example, one

system may generalise more than the other. In order to capture this key feature of our experimental paradigm in computational terms, we introduced a time-invariant weighting on the two systems that determines how much of what a system learns during prism exposure is subsequently expressed behaviourally on open-loop pointing. In practice the equation for CLP trials remains the same as the original model (equation 4.2). However, the equation for OLP becomes:

$$\begin{aligned}
 e(n) &= bias_{OLP} - X(n) \\
 X(n) &= w_f x_f(n) + w_s x_s(n) \\
 x_f(n) &= A_f x_f(n-1) \\
 x_s(n) &= A_s x_s(n-1) \\
 0 < A_f < A_s < 1; \quad 0 < B_s < B_f < 1 \\
 0 < w_f < 5; \quad 0 < w_s < 5
 \end{aligned} \tag{4.4}$$

where w_f and w_s are weighting factors that introduce a dissociation between the state of a given system on CLP and its contribution to the total amount of adaptation on OLP. As a result, the AE is now modelled as a weighted sum of the states of the two systems. For each system, having a weight lower than 1 means a scaling down of its contribution to the prism AE, whereas a weight greater than 1 means an over-expression.

In summary, the model variant proposed in this chapter is composed of three fixed parameters (d , $bias_{CLP}$ and $bias_{OLP}$), computed from the raw data the same way as for the original two-state model, and six free parameters (A_f , A_s , B_f , B_s , w_f and w_s) estimated via optimisation. The ability of the two models (original and weighted) to fit the behavioural data collected in this thesis is compared in section 4.3, after correcting for the benefit from greater complexity of the model variant

(i.e. six degrees of freedom instead of four originally). The two models were fitted using a least-mean-squared-error procedure implemented with the function *fmincon* in MATLAB (The MathWorks inc., version R2014b).

4.2.3 Data acquisition technique

Data were collected using either a 'manual' (experiments 1 and 4) or 'automated' (experiments 2 and 3) experimental set-up.

The manual set-up included a chinrest mounted on a pointing table. The pointing table had three fixed visual targets (2 lateral and 1 central target). At the end of each reaching movement an experimenter measured the pointing error using a ruler tapped underneath the targets, on the edge of the table so that it was invisible to participants. On OLP trials, a cardboard was used to block participants' vision after the onset of the movement. The main advantages of this set-up are its simplicity and portability. It can easily be used in any experimental population, including stroke patients.




Figure 4.2: Automated prism adaptation set-up. Participants sat in a chinrest, looking at a horizontal 32 inches touch screen through a liquid crystal (LC) shutter. The touch screen was used to present visual targets and record touch position. The LC shutter was used to control precisely participants' visual feedback on closed-loop and open-loop pointing trials. Vertical wood panels mounted on the three free edges of the table restricted individuals' vision to the experimental set-up only.

Datasets 2 and 3 were collected using an automated set-up that I designed and built (Figure 4.2). This set-up included a chinrest mounted on a pointing table with a 32-inch LED screen embedded in it. A liquid crystal shutter (Dispersion film, Liquid Crystal Technologies, USA) was positioned in front of the chinrest so as to control participants' visual feedback. On all pointing trials (OLP and CLP), the shutter turned opaque 500 ms after the registered touch in order to limit visual information to the centre-out movement only (and not the return movement). On OLP trials, the shutter turned opaque as soon as participants released a button situated on the pole of the chinrest. The shutter only turned clear again when individuals initiated the next trial by pressing the button again. Compared to the manual set-up, this automated system allows a more precise control over the timing of the experiment. It also enables a more accurate and reproducible measure of pointing errors, in a way that does not rely on an experimenter. This limits potential experimenter-related confounds. To date, this set-up has not yet been tested in a clinical population.

4.2.4 Experimental procedures

Four experimental datasets are used in this chapter. The first one was collected to address the three questions mentioned in introduction. The other three were collected for chapters 5, 6 and 8, and are used here to replicate some of the findings of experiment 1.

Experiment 1. This experiment was conducted in a group of 25 young healthy right-handed individuals (11 females, mean age = 29.5 years, SD = 5.5) using the manual PA set-up. The PA protocol described in Figure 4.1B (i.e. baseline CLP, baseline OLP, 6 blocks of prism exposure interleaved with 6 blocks of AE measurement) was used to measure the evolving dynamics of AE during the adaptation phase. We next assessed AE decay during active de-adaptation, to get an interference measure

of the strength of the AE memory trace. Immediately after the adaptation phase, prisms were removed and participants again made CLP movements to the left and right targets. As a consequence of PA, they initially made leftward errors and used visual feedback to correct their errors rightward to regain baseline accuracy (i.e. they de-adapted). In order to measure the dynamics of AE decay and spontaneous recovery during active de-adaptation, 6 blocks of OLP (Figure 4.3A; de-adaptation phase, dAE1-6: 15 OLP per block) were interleaved over 6 blocks of CLP (Figure 4.3A; de-adaptation phase, dE1-6: 10 CLP trials per block; equal number of trials to each of two target locations). The measure of the AE during active de-adaptation allowed cross-validation, i.e. fitting of the model on the adaptation phase and assessment of its ability to predict immediate de-adaptation data (section 4.3.1). Finally, retention of prism AE was subsequently measured 24 hours after the end of the de-adaptation period (45 OLP trials) to assess the models' ability to predict the AE memory trace at a late retention time point.




Figure 4.3: Experimental protocols of experiment 1. Participants were instructed to point at visual targets situated in front of them. Visual feedback of their terminal endpoint error was available on closed-loop pointing (CLP, in red) but not on open-loop pointing (OLP, in black). After a block of CLP and OLP baseline, participants alternated between blocks of prism exposure CLP (E1-2: 10 CLP trials per block; E3-6: 20 CLP trials per block) – during which they wore 10° right-shifting prism glasses and gradually corrected their pointing errors – and blocks of OLP (AE1-6: 15 OLP trials per block) – performed without prisms in order to measure the after-effect. The PA phase lasted about 20 minutes. Immediately after, participants underwent an active de-adaptation during which the same protocol was repeated without prisms (dE1-6: 10 CLP trials per block, dAE1-6: 15 OLP trials per block). Retention was probed on OLP 24 hours later (3 blocks of 15 trials).

All participants from experiments 2-4 underwent the same PA protocol (see section

4.2.1; Figure 4.1). Their PA data were used to compare the ability of the original and weighted two-state models to fit behaviour during PA (i.e. question 1 in introduction).

Experiment 2. This experiment was conducted in a group of 24 young healthy right-handed male individuals (0 female, mean age = 25.0 years, SD = 3.5) using the automated PA set-up. The full experimental protocol is detailed in chapter 5. All participants underwent the PA twice without brain stimulation so PA data averaged across these two sessions are used in this chapter.

Experiment 3. This experiment was conducted in a group of 23 elderly healthy right-handed male individuals (0 female, mean age = 67.2, SD = 8.9) using the automated PA set-up. The full experimental protocol is detailed in chapter 6.

Experiments 4. This experiment was conducted in a group of 15 spatial neglect patients (1 female, mean age = 59.2 years, SD = 7.8) using the manual PA set-up. Six of these patients participated in the rehabilitation study described in chapter 8. The PA data analysed in this chapter are taken from the first day of intervention with sham stimulation. Data from the other nine patients come from a different cohort, tested only once using the PA protocol described in section 4.2.1.

4.3 Results

4.3.1 Model comparison during the adaptation phase

Both the original and weighted two-state models were fitted to the adaptation phase of the group mean PA data of each experiment as described in section 4.2.2. The prism effect d was set to the group mean error on the first prism exposure trial for each dataset (dataset 1: $d = 7.06$; dataset 2: $d = 5.70$; dataset 3: $d = 6.06$; dataset 4: $d =$

3.33). Group fits are plotted in [Figure 4.4](#) and estimated parameters are summarised in [Table 4.1](#). In order to test the prediction that the weighted two-state model would explain our data better than the original two-state model (Smith *et al.* 2006a) while correcting for differences in model complexity, we performed formal model comparison using the Akaike’s Information Criterion (AIC), calculated as follows:

$$AIC = t \cdot \ln var(\epsilon) + 2k \quad (4.5)$$

where t is the total number of data points, k is the degrees of freedom in each model ($k = 4$ for the original model and $k = 6$ for the weighted model), and ϵ is the residual error of the model (i.e. the summation over all trials of the difference between the model prediction and observed data). This index expresses the residual error of the model while adding a penalty for model complexity in order to avoid issues related to over-fitting. Thus the model that yields the smallest AIC can be judged as the best model.

Across all four datasets, the weighted two-state model consistently displayed a closer fit to the group mean data (greater R-squared and lower root mean squared error, [Table 4.1](#)). The AIC confirmed that this improvement was not simply due to extra parameters (dataset 1: $\Delta AIC = -616.35$; dataset 2: $\Delta AIC = -569.92$; dataset 3: $\Delta AIC = -380.32$; dataset 4: $\Delta AIC = -156.96$).




Figure 4.4: Fit of the two models across the four datasets. The x-axis represents trial number and the y-axis pointing accuracy relative to baseline (lateral error, in cm, rightward errors are positive). Black wedges indicate blocks throughout which 10° right-shifting prisms are worn. Participants received visual feedback on CLP trials (shaded dark grey), allowing error-dependent learning to occur. By contrast, no visual error feedback was provided on OLP (shaded light grey), allowing the measurement of prism AE. Across all four datasets (A-D), the weighted two-state model (in red) yielded a closer fit to the data (group mean, error bar = SEM, in black) than the original model (in blue). Full details of the formal model comparison are shown in [Table 4.1](#).

The advantage of our variant over the original model was visible mainly at the beginning of blocks (Figure 4.4). The original model had a tendency to overshoot rightward on CLP trials and leftward on OLP trials compared to human behaviour. This was reflective of an excessively fast 'fast system' that learnt and decayed too quickly (Table 4.1: the original model yielded higher B_f and lower A_f). By contrast, the learning and forgetting dynamics of the slow system were comparable across the two models (Table 4.1: the original and weighted models yielded comparable B_s and A_s). By downscaling the contribution of the slow system to the AE (Table 4.1: $w_s < 1$ in experiments 1-3), the weighted two-state model allowed the fast system to learn and forget at a slower rate, which better fitted human behaviour.

	Exp. 1		Exp. 2		Exp. 3		Exp. 4	
	<i>Orig.</i>	<i>Weight.</i>	<i>Orig.</i>	<i>Weight.</i>	<i>Orig.</i>	<i>Weight.</i>	<i>Orig.</i>	<i>Weight.</i>
A_f	0.64	0.84	0.72	0.90	0.71	0.86	0.95	0.89
A_s	0.99	1.00	0.98	1.00	0.99	1.00	1.00	0.99
B_f	0.57	0.14	0.42	0.18	0.41	0.19	0.14	0.16
B_s	0.09	0.11	0.12	0.11	0.14	0.14	0.13	0.13
w_f		0.67		0.96		0.88		1.47
w_s		0.58		0.65		0.75		1.19
AIC	564.33	-52.02	314.60	-255.32	281.10	-99.22	132.67	-24.29
R^2	0.87	0.99	0.94	0.99	0.96	0.99	0.97	0.98
$RMSE$	0.75	0.21	0.45	0.14	0.42	0.19	0.31	0.22

Table 4.1: Model comparison for the four datasets. Across all datasets, the weighted model displayed lower Akaike information criterion (AIC), greater R-squared (R^2), and lower root mean squared error (RMSE), indicating a better fit to the data than the original two-state model.

Next, the two models were fitted to individual PA data to compare their performance at the individual level. The fitting procedure was identical to the one used at the group level. The prism effect d was calculated as the group mean shift on the first prism exposure trial and $bias_{CLP}$ and $bias_{OLP}$ were calculated as individuals' block mean error on the baseline CLP and OLP blocks, respectively. The results confirmed the superiority of the weighted two-state model at the individual level for experiments involving young healthy adults (Figure 4.5A-B; Wilcoxon rank sum test with

a Bonferroni significance threshold of $p = 0.013$; experiment 1: $p < 10^{-4}$, experiment 2: $p < 10^{-3}$). The improvement of individual fits induced by the weighted two-state model was close to significance level in the other two experiments (Figure 4.5C-D; Wilcoxon rank sum test with a Bonferroni significance threshold of $p = 0.013$; experiment 3: $p = 0.05$, experiment 4: $p = 0.16$).





Figure 4.5: Comparison of the Akaike's Information Criterion among the two models for individual fits. Box plots of the AIC calculated on the fit of individual PA datasets for the original (in blue) and weighted (in red) two-state model across all experiments (A-D). After Bonferroni correction (threshold $p < 0.0125$), the weighted model significantly outperformed the original model at the individual level in young healthy adults (A, B; both $p < 10^{-3}$) and there was a trend in elderly adults (C; $p = 0.05$) and neglect patients (D; $p = 0.16$). On each box, the central mark indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme data points not considering outliers, which are represented as the '+' symbol. Grey lines represent each individual.

4.3.2 Predictive power during active de-adaptation

In experiment 1, participants underwent an active de-adaptation immediately following PA to provide an interference measure of the strength of the AE memory trace. The de-adaptation phase was used to compare the models' predictive power, using a cross-validation approach. First, the original and weighted two-state models were fitted to the group mean data during the adaptation phase (Figure 4.6, phase in green). Second, predicted de-adaptation data were generated based on the model parameters estimated during the adaptation phase (Figure 4.6, phase in black). The comparison of the real versus predicted de-adaptation data enabled us to evaluate the two models' ability to predict future motor behaviour after having been trained on a subset of the

data (cross-validation).

A) Original two-state model



B) Weighted two-state model




Figure 4.6: Descriptive and predictive value of the weighted two-state model. The x-axis represents trial number and the y-axis pointing accuracy in reference to baseline accuracy (lateral error, in cm, rightward errors are positive). Black wedges indicate blocks throughout which 10° right-shifting prisms are worn. Participants received visual feedback of their reach endpoint error on CLP trials (shaded dark grey), but not on OLP (shaded light grey). Cross-validation of the two models was performed by fitting the model to the adaptation phase (green) and comparing the two models' ability to predict immediate de-adaptation performance (black). The weighted model (B) outperformed the original model (A) ($\Delta AIC = -166.86$).

In order to control for model complexity, we used the AIC – calculated on the de-adaptation phase only – as our metric of predictive power. As expected, the weighted two-state model showed better predictive power than the original model, as expressed

by a lower AIC (original model: $AIC = 232.47$; weighted model: $AIC = 66.61$, Figure 4.6). This advantage was particularly evident early during de-adaptation (Figure 4.6: dE1-2, dAE1-2), when the two systems were still in an adapted state.

4.3.3 Predictive power 24 hours after the adaptation

So far, the results have demonstrated that the introduction of a weighted sum on open-loop pointing trials allowed the two-state model to better fit human data during PA, and better predict subsequent behaviour in the period immediately following adaptation (active de-adaptation). Here, we asked whether the weighted two-state model could also predict the AE memory trace 24 hours later. From a conceptual perspective, the slow system is associated with memory persistence (i.e. long memory timescale), whereas the fast system is associated with rapid memory decay (i.e. short memory timescale) (Körding *et al.* 2007; Joiner & Smith 2008; Inoue *et al.* 2015). Thus, we hypothesised that a greater contribution of the slow system to the AE during the adaptation should predict greater and/or more stable AE on the next day. The weighted two-state model provides a time-invariant estimate of the contribution of the fast and slow systems to the AE during PA, namely the weighting parameters w_f and w_s . We could therefore derive a relative weight parameter R_w as follows:

$$R_w = \frac{w_s}{w_f + w_s} \quad (4.6)$$

This derived scalar captures in a single metric the relative contribution of long versus short memory timescales to the AE as they develop during the adaptation. A relative weight greater than 0.5 is indicative of a dominance of the slow system in OLP behaviour. This entails little intra- and inter- block forgetting during the adaptation. Alternatively a fast system dominated PA profile ($R_w < 0.5$) is characterised by large prism AE initially, but greater forgetting throughout the 20 minutes of PA.




Figure 4.7: After-effect memory trace 24 hours after prism adaptation (Experiment 1). The x-axis represents trial number (within the 24-hours retention probe) and the y-axis pointing accuracy in reference to baseline accuracy (lateral error, in cm, rightward errors are positive). Each individual is represented as an individual line, colour-coded by the relative weight R_w (estimated on the adaptation data, collected 24 hours earlier). The black dots represent, for each trial, the group mean after-effect.

A long retention probe of 45 OLP trials was used in experiment 1 to measure the AE memory trace 24 hours after the adaptation. [Figure 4.7](#) shows individuals' behavioural data during this block. No significant AE was detected at the group level during this block (mean AE = -0.35, SEM = 0.26, $t_{(24)} = -1.35$, $p = 0.19$; [Figure 4.1](#)). This is likely to be the result of the de-adaptation performed the day before ([Figure 4.6](#)). Nevertheless, the observation of individual datasets revealed certain dynamics within this block that may reflect memory reactivation as participants were re-exposed to the experimental set-up ([Figure 4.7](#)). To address this question, the magnitude of prism AE was calculated separately for the first and last 15 trials of this long exposure block so as to obtain a distinct measure of 'early' and 'late' AE memory trace at 24 hours. These two measures were entered as independent variables in a robust multiple linear regression that included the relative weight R_w as the dependent variable. Both the early and late AEs correlated independently and inversely with the relative weight (regression adjusted R-squared = 0.33; early: $\beta = 0.93, p = 0.001$; late: $\beta = -0.57, p = 0.03$, [Figure 4.8](#)). This result indicates

that, across individuals, greater prism AE at the early phase of the 24-hours retention probe was associated with a greater contribution of the fast system to the AE during the adaptation phase (on the previous day). Alternatively, greater prism AE at the later phase of the retention block was associated with a greater contribution of the slow system to the AE during the adaptation.




Figure 4.8: The dynamics of long-term memory reactivation reflect the relative weighting between fast and slow systems. Individuals with greater relative weight of the slow system during adaptation had lower prism AE at the early phase (left plot) but greater prism AE at the later phase (right plot) of the 24-hours retention block. The colour code represents the robust multiple linear weight. In red is the regression fit with 95% confidence interval.

The same analysis was run on the baseline OLP block of day 1 (also 45 trials) to verify that these relationships were not confounded by natural directional shift in pointing accuracy occurring regardless of the adaptation. None of the early or late baseline OLP errors were found to correlate with the relative weight (both $p > 0.2$).

4.4 Discussion

This experimental chapter aimed to develop a computational framework within which to describe and explain the AE memory trace. We proposed a variant of the linear time-invariant two-state model (Smith *et al.* 2006a) that enables a scaling of the

contribution of the fast and slow systems to open-loop pointing behaviour. This modification allowed the generation of a range of possible AE dynamics based on a linear combination of the memory signals generated during prism exposure. Across all PA experiments reported in this thesis (young healthy adults, elderly healthy adults, neglect patients), the weighted two-state model outperformed the original model during the adaptation phase. Our variant was also better at predicting de-adaptation data immediately following PA. Finally, we found evidence that, across individuals, greater relative contribution of the slow system to the AE during the adaptation is associated with slower but greater memory reactivation when individuals are tested again 24 hours later.

The framework presented in this chapter provides a clear mathematical description of putative internal states driving PA and after-effects. It constitutes a clear advantage over the traditional empirical approach because it offers a means to generate and test quantitative behavioural predictions (Petitet *et al.* 2017). For example, based on the cross-sectional relationships found in this chapter, one may predict that increasing the relative contribution of the slow system to the after-effect would result in a more stable memory trace.

The key conceptual advance of our model is the idea that the behavioural expression of the fast and slow systems varies across individuals, and depending on the task constraints. Compared to the original model (Smith *et al.* 2006a), this adds another degree of freedom to explain inter-individual differences. Indeed, human participants may not only differ in the speed at which their fast and slow systems learn and forget, but also in terms of the degree to which these systems transfer across task constraints. By only considering one AE measure (i.e. OLP to a central target), our experimental paradigm provides the simplest measure of generalisation, which is an advantage in

clinical practice (Jacquin-Courtois *et al.* 2013). Hence, it only required the introduction of one weighting parameter per system (w_f and w_s). However, an interesting follow-up question is to determine the dependency of these weights on task parameters such as position of the target, speed of the reaching movement or limb used. These are important experimental parameters known to influence the magnitude of the AE in healthy individuals (Hamilton & Bossom 1964; Bedford 1989; Kitazawa *et al.* 1997; Michel *et al.* 2007). The computational framework we proposed offers a way of quantifying generalisation of the two systems to various contexts and mathematically describe their transfer functions.

The following chapters build on the conceptual framework developed in this chapter to investigate the neuro-computational mechanisms underlying memory persistence following PA.

Chapter 5

Causal explanation of individual differences in human sensorimotor memory formation

Previous work from the laboratory demonstrated that excitatory stimulation of the sensorimotor cortex during prism adaptation enhances retention of the after-effect in young healthy adults (O'Shea et al. 2017). In this experimental chapter, we replicate this effect and characterise its causal neuro-computational mechanism, using multi-modal ultra high-field MRI and the computational model developed in the previous chapter.

5.1 Introduction

Individuals differ – an explanatory challenge. Here we exploit this variation to test a causal mechanistic hypothesis about how sensorimotor memories in the human brain are formed, retained and enhanced. Primary motor cortex (M1) has been identified as a critical region for the maintenance of newly acquired visuomotor maps following sensorimotor adaptation in both primates (Wise *et al.* 1998; Li *et al.* 2001; Paz *et al.*

2003; Paz *et al.* 2005; Inoue *et al.* 2016) and humans (Hadipour-Niktarash *et al.* 2007; Hunter *et al.* 2009; Galea *et al.* 2010; Landi *et al.* 2011; Leow *et al.* 2016; O’Shea *et al.* 2017). Previous work from our laboratory demonstrated that left M1 anodal transcranial direct current stimulation (a-tDCS) applied during adaptation to a 10° rightward displacement of the visual field (prism adaptation, PA) enhances consolidation of the prism after-effect (AE) in the healthy brain (O’Shea *et al.* 2017). Notably, this stimulation-induced memory enhancement was 1) *anatomically specific* (it was not observed with left parietal or right cerebellar tDCS), 2) *polarity specific* (it was not observed with left M1 cathodal tDCS), and 3) *cognitive state dependent* (it was not observed when a-tDCS was applied before as opposed to during PA). Here, we interrogate the computational and neural bases of this effect.

Since adaptation memory was enhanced only with M1 excitatory stimulation during PA, the interaction with information processing in M1 was critical. What computation did stimulation alter to cause memory change? Computational theories of motor control posit that when a perturbation (such as an optical shift) causes reaching errors, the brain forms an estimate of the perturbation, used to counteract the disturbance and correct movement accuracy (Shadmehr *et al.* 2010; Franklin & Wolpert 2011; Wolpert *et al.* 2011; Petitet *et al.* 2017). For future movements to also be accurate, this estimate should not only include a representation of the magnitude of the perturbation (e.g. by how much do I need to correct?) but also of its temporal dynamics (e.g. for how long is the perturbation likely to prevail?). The latter is known as a *temporal credit assignment* problem. For retention to be optimal, it should match the internal estimate of the perturbation timescale, i.e. the solution to the temporal credit assignment (Körding *et al.* 2007). For example, if the disturbance is believed to be stable, adaptation should be retained; if transient, it should decay. Here we reasoned that it might be this computation, whose implementation is likely

to involve M1, that interacts with a-tDCS such that memory is enhanced.

If that is true, what features of M1 physiology are likely to be involved in this computation? Primary motor cortex a-tDCS is known to locally increase the excitation:inhibition (E:I) ratio via a reduction of gamma-aminobutyric acid (GABA) concentration (Stagg *et al.* 2009; Kim *et al.* 2014; Bachtiar *et al.* 2015; Antonenko *et al.* 2017). In our previous study, the magnitude of the stimulation-induced GABA decrease (measured in a separate session, inside the MRI scanner) correlated with the increase in AE 24 hours after receiving PA + left M1 a-tDCS (O'Shea *et al.* 2017), implicating this neurotransmitter in the mechanisms responsible for the memory enhancement. Therefore, we hypothesised that the solution to the temporal credit assignment might be reflected in the level of M1 excitation:inhibition during PA, such that greater E:I ratio would be associated with longer memory timescales. We used the weighted two-state model described in the previous chapter ([chapter 4](#)) to infer this computation (i.e. relative contribution of short-lasting versus long-lasting computational states to the overall representation of the perturbation) at the individual level, as participants adapted to prisms while or after receiving real or sham left M1 a-tDCS. Individual trait markers of neurochemical response to M1 a-tDCS, white matter microstructure, resting functional connectivity, as well as BDNF val66met polymorphism were taken to test the main prediction and explain individual variability in computation and behaviour. The data presented in this chapter demonstrate that sensorimotor cortex E:I (and in particular GABAergic inhibition) causally determines the time constant of memory persistence, explaining phenotypic variation in adaptation decay in young healthy adults.

5.2 Material and Methods

5.2.1 Participants

Twenty-four right-handed healthy male individuals (mean age = 25.03 years, SD = 3.50) gave their written informed consent to take part in this study in accordance with ethical approval from the Oxford A Research Ethics Committee (REC reference number: 13/SC/0163). The gender selection criterion was applied to avoid potential confounds related to variations in neurotransmitter concentration with the menstrual cycle in women (Smith *et al.* 1999; Epperson *et al.* 2002). Handedness was assessed using the Edinburgh handedness questionnaire (Oldfield 1971). All participants had normal or corrected-to-normal vision and indicated no family history of psychiatric or neurological disease. All participants were naive to the purpose of the experiment and were not informed of the expected effect of the prism glasses and brain stimulation. No participant reported any side effect from the experimental procedure.

5.2.2 Experimental design

All individuals participated in four PA sessions (behavioural experiment) and two scanning sessions (neuroimaging experiment). Sessions were separated by at least one week. The order of the four PA sessions was fully counterbalanced across the group. The order of the scanning sessions was also counterbalanced across the group.

The behavioural experiment had a 2-by-2 within-subject design. Experimental sessions differed in regards to the type (anodal or sham tDCS) and timing (during or before PA, i.e. online or offline) of the stimulation. Similar to our previous report (O'Shea *et al.* 2017), each PA session included four phases as follow: 1) measure of baseline pointing accuracy (3 blocks); 2) prism adaptation (14 blocks; 20 min); 4) 10-minute break (blindfolded, at rest); 3) measure of prism AE retention ([Figure 5.1B](#)).

The neuroimaging experiment (performed in the same individuals) involved 2 scanning sessions with either real or sham left M1 tDCS. Details of the scanning protocol is provided in [Figure 5.1C](#).




Figure 5.1: Experimental set-up and protocols. **A.** The automated PA set-up was used (see section 4.2.3). **B-C)** Experimental protocol. Closed-loop pointing (CLP) and open-loop pointing (OLP) accuracy was first measured at baseline. Then followed 20 minutes of either real or sham left M1 a-tDCS during which participants either adapted to 10-degree right-shifting prisms (online condition, B) or stayed blindfolded at rest (offline condition, C). In the offline condition, participants underwent prism adaptation as soon as the stimulation finished. The prism adaptation phase included blocks of CLP during which participants adapted to prisms, interleaved with blocks of OLP used to measure the development of prism AE. After the end of the last block of the PA phase, participants rested blindfolded for 10 minutes before to undergo a long block of 45 OLP trials to measure retention. **D.** Scanning protocol. All participants underwent two scanning sessions with either anodal or sham left M1tDCS, the order of which was counterbalanced across the group. MR spectroscopy and resting state fMRI were acquired at two time points: before and after tDCS. Diffusion-weighted images were taken once at the end of one of the two scanning sessions in 22 subjects. Timeline shown is an estimate of the length of the scans in minutes.

5.2.3 Prism adaptation procedure

The experimental protocol used in this study was as described in the previous chapter ([chapter 4](#)) and is summarised below. The main difference with previous work

(O'Shea *et al.* 2017) is that a fully automated experimental set-up was used to enable a more precise and reproducible measurement of reaching errors, as well as a better control of visual feedback availability and experimental timing (Figure 5.1A). As a result, the pace of the experiment increased so we could fit 7 exposure blocks instead of the original 6 during the 20 minutes of brain stimulation (Figure 5.1BC).

Participants sat facing a horizontal 32-inch LED screen embedded in a table that was used to record the reach endpoint position of their index finger (Figure 5.1A). Their head was restrained by a chinrest mounted on the edge of the table to maintain a distance of about 60 centimetres between their eyes and the centre of the screen. Wood panels were placed on the other three edges of the table to create a box in which participants' visual environment was strictly restricted to the screen. A button was attached to the pole of the chinrest and subjects were instructed to keep it pressed at all times and only release it when initiating a reaching movement. A fixed shutter prevented participants from seeing their limb when pressing the button and during the first third of their pointing movement. A liquid crystal shutter (Dispersion film, Liquid Crystal Technologies, Ohio, USA), placed between participants' eyes and their limb, controlled the visual feedback of the screen and pointing limb. The same prism glasses were used (glacier goggles: Julbo, Longchaumois, France; 10° right-shifting prism lenses: OptiquePeter, Lyon, France).

Participants were instructed to perform two types of pointing movements: closed-loop pointing and open-loop pointing.

- **Closed-loop Pointing (LP).** On CLP trials, participants made speeded reaching movements (average movement duration = 317 ms, SD = 53 ms) to point at a visual target appearing on the screen at the beginning of each trial. Within any CLP block, targets appeared either 12 cm to the left or 12 cm to the right

of the centre of the screen (50% on each side), following a pseudorandom sequence in which target location transitions were generated if the same location appeared twice in a row. Visual feedback was limited to the last two thirds of the reaching movement in order to limit strategic adjustment (Redding & Wallace 1996; O'Shea *et al.* 2014; Inoue *et al.* 2015; O'Shea *et al.* 2017) and lasted until 500 ms after participants touched the screen to allow for them to perceive their terminal reach endpoint error. After this time, the LC shutter turned opaque and participants had to return to the starting position (i.e. press and hold the button) without visual feedback of their hand.

- **Open-loop pointing (OLP).** On OLP trials, participants were instructed to point at a slower speed (average movement duration = 816 ms, SD = 174 ms), to a visual target located in the centre of the screen. The LC shutter turned opaque as soon as their hand left the starting position (i.e. release of the button), thereby blocking any visual feedback during the reaching and return movements. The LC shutter turned clear again only when participants returned to the starting position and initiated the next trial.

Four factors differentiated OLP from CLP trials: 1) participants only wore prism glasses during CLP trials (and never during OLP trials) so that any change in pointing accuracy on open-loop pointing relative to baseline could be interpreted as an after-effect of the adaptation; 2) participants could learn from their reach endpoint error on CLP but not on OLP (because visual feedback was deprived); 3) participants made fast movements on CLP (average movement duration = 317 ms, SD = 53 ms) but slower movements on OLP (average movement duration = 816 ms, SD = 174 ms); 4) participants pointed to lateral targets on CLP (located 12 cm to the left or right of the centre of the screen) and to a unique central target on OLP. The differences in movement speed and target location between prism exposure (CLP trials) and AE probes (OLP) enabled us to assay generalised AE that were not contaminated by ex-

tensive training of a specific movement dynamic and local learning at a certain target location (Kitazawa *et al.* 1997). This type of AE is thought to be relevant for neglect rehabilitation (Serino *et al.* 2006; O’Shea *et al.* 2017).

In all four PA sessions, the tDCS electrodes were positioned and the stimulation switched on at the end of the baseline phase (20 CLPs, 30 OLPs). In the online condition, the PA phase (E1-7: 10 CLPs per block during the first 2 blocks, 20 CLPs per block in the remaining ones; AE1-7: 15 OLPs per block) started as soon as the stimulation reached its constant level. In the offline condition, participants were blindfolded and remained at rest during the entire duration of the stimulation. This enabled us to control the visual input participants received during the offline stimulation. The prism adaptation phase started as soon as the stimulation was over in the offline condition. In all sessions, retention of the AE (AE8-10: 15 trials per block) was probed 10 minutes after the end of the last OLP trial of the adaptation phase.

5.2.4 Transcranial direct current stimulation

5.2.4.1 During behavioural sessions

When delivered outside the scanner (i.e. PA sessions), direct current stimulation was generated by a battery driven DC stimulator (Neuroconn GmbH, Ilmenau, Germany) connected to two 7×5 cm sponge electrodes soaked in a 0.9% saline solution. The anode electrode was centred over C3 (5 cm lateral to Cz) corresponding to the left primary motor cortex according to the international 10-20 System (Herwig *et al.* 2003). The cathode electrode was placed over the right supraorbital ridge. The large surface of the electrode (35 cm²) ensured a good coverage of the M1 “hotspot”. Similar to previous M1 tDCS studies, the long axis of the anode electrode was oriented parallel to the sagittal axis during behavioural sessions (Nitsche & Paulus 2000; Galea *et al.*

2010; Panouillères *et al.* 2015; O'Shea *et al.* 2017). The electrodes were positioned immediately before stimulation onset and removed as soon as the stimulation finished in order to minimise participants' discomfort. For anodal stimulation, the current intensity was set to 1 mA for 20 minutes with a ramp-up and ramp-down period of 10 seconds prior to and after the stimulation respectively. For sham stimulation, the current also ramped up and down for 10 secs but no stimulation was delivered during the 20 minutes. Instead, small current pulses (110 μ A over 15 ms) occurred every 550 ms to simulate the tingling sensation associated with real anodal stimulation. Both experimenters and participants were blinded to the stimulation condition. In order to verify the efficacy of the blinding, participants were asked to guess whether the stimulation was 'real' or 'fake' at the end of every behavioural session. The proportion of correct responses for identifying the stimulation condition was at chance level (50%, SD = 27.59), indicating that participants were unable to successfully differentiate anodal from sham tDCS.

5.2.4.2 Inside the MRI scanner

When delivered inside the MRI scanner (i.e. PA sessions), direct current stimulation was generated by a battery driven DC stimulator (Neuroconn GmbH, Ilmenau, Germany) connected to two 7×5 cm MR-compatible carbon-rubber electrodes equipped with an inbuilt 5 k Ω resistor at the junction between the electrode and the wire (Easy-cap GmbH, Herrsching, Germany) in order to avoid temperature increase and reduce induction voltages caused by the scanner (Woods *et al.* 2016). A thick coating of chloride-free abrasive electrolyte gel was applied between the biocarbon electrodes and the scalp to maintain low levels of impedance throughout the duration of the scan and ensure the stimulation to be delivered evenly across the electrode (Woods *et al.* 2016). The electrode placement was identical except the anode electrode was oriented perpendicular to the sagittal axis (instead of parallel, during behavioural

sessions) in order for the wire to exit the RF coil through the shortest route and limit the risk of eddy current induction (Woods *et al.* 2016). This electrode montage was similar to previous work done in the laboratory (Stagg *et al.* 2009; Bachtiar *et al.* 2015; O'Shea *et al.* 2017). The electrodes were placed on participants' head immediately before the beginning of the scanning session and removed as soon as participants exited the scanner. The protocol for anodal and sham stimulation was identical to that delivered outside the MRI scanner (i.e. 1 mA, 20 minutes). Only participants were blinded to the stimulation condition inside the scanner.

5.2.5 Statistical analysis of raw prism adaptation data

Reach endpoint angular errors were calculated as the angle formed between a straight line joining the starting position (button) and the target location and a straight line joining the starting position (button) and the touch location on the screen. Data distribution assumption of normality was assessed quantitatively using the Kolmogorov-Smirnov and Shapiro-Wilkinson tests and qualitatively using distribution plots and Q-Q plots. Sphericity was assessed using Mauchly's test, with violations corrected using the Huyhn-Feldt procedure where appropriate. When data met the assumptions of parametric testing, pointing angular error data were analysed with repeated measures ANOVA and t-tests using SPSS software (SPSS inc. version 24.0) and Matlab (The MathWorks inc. version R2014b). Cohen's d was reported for significant results only.

5.2.6 Computational analysis of prism adaptation data

The weighted two-state linear time-invariant model described in [chapter 4](#) was fitted to the adaptation data. On CLP trials, the model is defined as follows:

$$\begin{aligned}
 e(n) &= bias_{CLP} + d - X(n) \\
 X(n) &= x_f(n) + x_s(n) \\
 x_f(n) &= A_f x_f(n-1) + B_f e(n-1) \\
 x_s(n) &= A_s x_s(n-1) + B_s e(n-1) \\
 0 < A_f < A_s < 1; \quad 0 < B_s < B_f < 1
 \end{aligned} \tag{5.1}$$

where the endpoint error on the n -th trial $e(n)$ is modelled as the sum of a bias term $bias_{CLP}$ and a prism effect d from which is removed the sum of the states of the two systems $X(n)$. The learning and retention rates (B and A) determine the speed of learning and decay of each system. On OLP trials, where prisms are removed and no visual feedback is available, the weighted two-state model becomes:

$$\begin{aligned}
 e(n) &= bias_{OLP} - X(n) \\
 X(n) &= w_f x_f(n) + w_s x_s(n) \\
 x_f(n) &= A_f x_f(n-1) \\
 x_s(n) &= A_s x_s(n-1) \\
 0 < A_f < A_s < 1; \quad 0 < B_s < B_f < 1 \\
 0 < w_f < 5; \quad 0 < w_s < 5
 \end{aligned} \tag{5.2}$$

Equations [6.1](#) and [8.6](#) differ in four ways. First, a different bias term is used because of baseline differences in pointing accuracy between CLP and OLP. Second, the prism effect is set to zero to reflect removal of the prisms. Third, the learning rates of both systems (B_f , B_s) are set to zero because the absence of visual feedback prevents

error-dependent learning on OLP. Fourth, weighting factors (w_f, w_s) are introduced to allow for a differential contribution of the fast and slow systems to motor behaviour during AE measurements (see [chapter 4](#)).

Model optimisation was performed in MATLAB (The MathWorks inc., version R2014b) using a least-mean-squared-error fitting procedure (*fmincon* function). In order to confirm the superiority of the weighted two-state model over the original non-weighted two-state model (Smith *et al.* 2006a), we performed formal model comparison on the group mean data (average of the four sessions of the 24 participants, i.e. 96 datasets). For the original two-state model, the best-fitted parameters were $A_f = 0.866$; $A_s = 0.991$; $B_f = 0.195$ and $B_s = 0.101$ and for the weighted two-state model, they were $A_f = 0.931$; $A_s = 0.996$; $B_f = 0.159$; $B_s = 0.076$; $w_f = 0.974$ and $w_s = 0.855$ (Supplementary scripts). The overall fit of the group mean data was better for the weighted two-state model compared to the original model (root mean squared error = 0.23° vs. 15.48° ; R-square = 99.68% vs. 99.14%). In order to formally compare the two models, Akaike's Information Criterion (AIC) and the associated relative likelihood (RL) were calculated for each model as follow:

$$\begin{aligned} AIC &= t \cdot \ln \text{var}(\epsilon) + 2k \\ RL &= \exp\left(\frac{\Delta AIC}{2}\right) \end{aligned} \tag{5.3}$$

where t is the total number of data points and k is the degrees of freedom in each model ($k = 4$ for the original model and $k = 6$ for the weighted model). This index expresses the residual error while adding a penalty for model complexity in order to avoid issues related to over-fitting. Thus, the best model is characterised by the lowest AIC (i.e. a negative ΔAIC). This analysis confirmed that the improvement associated with the weighted two-state model was not merely due to extra parameters ($\Delta AIC = -273.02$, $RL = 3.06 \times 10^{-61}$).

The main advantage of the weighted two-state model is not only that it provides a closer fit to the data, but mainly that it enables us to extract the key computational parameter of interest for this study, i.e. the time-invariant contribution of the fast and slow systems to overt behaviour on OLP during adaptation. We expressed this as a relative weight as follow:

$$R_w = \frac{w_s}{w_f + w_s} \quad (5.4)$$

This derived scalar captures in a single metric the relative contribution of long versus short memory timescales to the AE during adaptation. A relative weight greater than 0.5 is indicative of a dominance of the slow system, i.e. little intra- and inter- block forgetting (Figure 5.4C). Alternatively a fast system dominant PA profile ($R_w < 0.5$) is characterised by larger prism AE initially but greater forgetting throughout the 20 minutes of prism adaptation (Figure 5.4C).

In order to ensure comparability of the relative weights within and between participants, A_f , B_f , A_s and B_s were kept constant and only w_f and w_s were estimated when fitting individual datasets. We used the learning and retention factors extracted from the fit of the weighted two-state model to the group mean data. If all parameters were let free, weighting factors would refer to different memory timescales depending on which participant/session they apply to. Instead, we decided to posit identical systems for all participants and only let the weight of these systems vary. This approach also reduced the number of free parameters, making the estimates more robust and ensuring that they captured unique variance. Furthermore, inspection of raw PA data revealed a stereotypical behavioural profile on CLP while most of the variance was present in the OLP data (Figure 5.4B).

5.2.7 Regression analyses

Regression analyses were performed in MATLAB (The MathWorks inc. version R2014b) using robust multiple linear regressions analyses with an iteratively re-weighted least squares procedure. This algorithm was used to limit the contribution of potential outliers to the fit. Individuals with missing data for at least one of the variables included in the linear regression were excluded from the analysis. The details of participants exclusion is provided in [Table 5.1](#). For all multiple linear regressions, normal probability plots of residuals were examined to check the normal distribution of errors. The coefficients reported, noted β , are the regression coefficient estimates after z-transformation of the dependent and independent variables.

5.2.8 BDNF Val66Met polymorphism

A sub-sample of 21 participants ([Table 5.1](#)) was genotyped for the BDNF Val66Met polymorphism (rs6265, G > A). Genomic DNA was extracted from buccal cells using the ChargeSwitch® gDNA Buccal Cell Kit (ThermoFisher Scientific, UK) and samples were genotyped in duplicate by LGC Genomics (LGC Group, UK). Rs6265 was the only polymorphism examined.

5.2.9 MR data acquisition

All MR data (T1-weighted images, MR spectroscopy, diffusion-weighted images and resting state fMRI) were acquired on a 7T Siemens MAGNETOM scanner (Siemens, Erlangen, Germany) at the Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB). All participants underwent two 90 minutes scanning sessions (anodal tDCS scan and sham tDCS scan) separated by at least one week, and counterbalanced across participants ([Figure 5.1D](#)). To keep volunteers engaged during the scanning session, a video (BBC life, no sound, no subtitles) was projected

in the scanner, except during the resting state fMRI acquisitions. Participants wore earplugs throughout scanning.

5.2.9.1 T1-weighted structural images

High-resolution T1 images (MPRAGE: 176×1 mm sagittal slices; TR/TE = 2200/2.82 ms; flip angle = 7 deg; FOV = 256×256 ; voxel size = 1mm isotropic; PAT = 2; scan time = 325 secs) were acquired at the beginning of each scanning session for MRS voxel placement and registration purposes.

<i>Participant</i>	GABA				Glu				tCr				DTI	rs-fMRI	BDNF
1	■									■					
2															
3															
4		■				■				■					
5															
6		■	■												
7									■						
8												■			
9															
10												■	■	■	
11													■	■	
12															
13		■									■				
14															
15		■													
16															
17			■												
18															■
19															
20															■
21															■
22			■												
23															
24															

Table 5.1: Diagram of data completeness. Completeness of the dataset is represented for each participant. The four columns of GABA, Glutamate and Creatine represent the four measurements (before a-tDCS, after a-tDCS, before s-tDCS, after s-tDCS). A black square means that the data didn't meet the quality criteria. Black squares in the BDNF column indicates participants for which genotyping could not be obtained.

5.2.9.2 MR spectroscopy

MR spectroscopy was acquired immediately before and after the application of 20 minutes of left M1 real or sham a-tDCS inside the scanner. B0 shimming was performed in a two-step process. First, GRE-SHIM ($64 \times 4\text{mm}$ axial slices, $\text{TR} = 600\text{ ms}$; $\text{TE}_{1/2} = 2.04/4.08\text{ ms}$; flip angle = 15° ; $\text{FOV} = 384 \times 384$; voxel size = 4 mm isotropic; scan time = 44 secs) was used to determine the optimal first- and second-order shim current. The second step involved the fine adjustment of first-order shims using FASTMAP. Prior to switching on the direct current stimulator, the modified semi-LASER sequence (O'Shea *et al.* 2007; Bank *et al.* 2015) ($\text{TR}/\text{TE} = 5000/36\text{ ms}$, 64 scan averages, scan time = 320 secs) was used to acquire MRS measurement in a $2 \times 2 \times 2\text{ cm}^3$ volume of interest centred on the left motor knob (Yousry *et al.*, 1997) to include parts of the pre- and post-central gyrus (see VOI overlap maps in Figure 5.10B). As soon as the baseline MRS measurement was over, the experimenter turned on the DC stimulator and no scanning happened for 10 minutes (Figure 5.1D). At the end of this time, lower-resolution T1-weighted images (MPRAGE: $176 \times 1\text{ mm}$ sagittal slices; $\text{TR}/\text{TE} = 2200/2.82\text{ ms}$; flip angle = 7° ; $\text{FOV} = 192 \times 192$; voxel size = 1 mm isotropic; $\text{PAT} = 4$; scan time = 171 secs) were acquired to check for head motion. If significant head motion was detected, the VOI was re-positioned to match the baseline acquisition and B0 shimming was performed again. If no significant head motion was detected, baseline VOI positioning and shimming were applied. On average, post-tDCS MRS measurements (64 averages) started 2 minutes 21 seconds after the end of the stimulation (range: 6 secs to 12 minutes 30 secs).

5.2.9.3 Diffusion weighted images (DWI)

Diffusion-weighted echo planar images (2 acquisitions of 64 diffusion-weighted directions in opposite phase-encode direction, $b\text{-value} = 1000\text{ s/mm}^2$; 8 non-diffusion weighted images with four b_0 volumes in the reverse phase-encode direction, $b\text{-value}$

= 0 s/mm²; 120 slices; TR/TE = 6210/68.2 ms; FOV = 192 × 192; voxel size = 1.2 mm isotropic; PAT = 2; scan time = 938 secs) were acquired once at the end of a scanning session for 22 participants (Figure 5.1D; Table 5.1).

5.2.9.4 Resting state functional MRI (rs-fMRI)

Resting state Blood oxygen-level-dependent (BOLD) echo planar fMRI scans (96 slices; TR/TE = 1472/25.0 ms; FOV = 192 × 192; voxel size = 1.5 mm isotropic; PAT = 2; scan time = 444 secs) were acquired before and after the stimulation (Figure 5.1D). Participants were instructed to keep their eyes open and fixate on a cross hair projected on the screen. Immediately after each resting state fMRI scan, a field map (64 slices; TR = 620 ms; TE1/2 = 4.08/5.1 ms; FOV = 192 × 192; voxel size = 2 mm isotropic; scan time = 122 secs) was acquired for post-processing correction of geometrical distortions caused by magnetic field inhomogeneity. Two participants did not undergo the resting state scans in order to minimise the time they spent in the scanner (Table 5.1).

5.2.10 MR data analysis

5.2.10.1 MR spectroscopy

Metabolites were quantified using LCModel (Provencher 1993; Provencher 2001; Provencher 2012) performed on all spectra within the chemical shift range 0.5 to 4.2 ppm. The model spectra of alanine (Ala), ascorbate/vitamin C (Asc), aspartate (Asp), glycerophosphocholine (GPC), phosphocholine (PCho), creatine (Cr), phosphocreatine (PCr), GABA, glucose (Glc), glutamine (Gln), glutamate (Glu), glutathione (GSH), myo-inositol (myo-Ins), Lactate, N-acetylaspartate (NAA), N-acetylaspartylglutamate (NAAG), phosphoethanolamine (PE), scyllo-inositol (scyllo-Ins) and taurine (Tau) were generated based on previously reported chemical shifts and coupling constants by VeSPA Project (Versatile Stimulation, Pulses and Analysis). The unsup-

pressed water signal acquired from the volume of interest was used to remove eddy current effects and to reconstruct the phased array spectra (Natt *et al.* 2005). Single scan spectra were corrected for frequency and phase variations induced by subject motion before summation. The full width half maximum (FWHM) and signal-to-noise ratio (SNR) did not differ between the four MRS measurements (Table 5.2).

	Anodal scan		Sham scan	
	<i>Baseline</i>	<i>post-tDCS</i>	<i>Baseline</i>	<i>post-tDCS</i>
% Grey matter	28.3 (4.7)	28.03 (5.2)	28.7 (3.2)	28.0 (3.7)
% White matter	65.3 (6.1)	65.5 (7.1)	64.4 (4.6)	65.4 (5.0)
% CSF	6.4 (2.5)	6.4 (2.7)	6.9 (2.3)	6.5 (2.4)
SNR	48.0 (5.3)	47.8 (6.6)	46.3 (7.3)	47.0 (5.8)
FWHM	0.028 (0.003)	0.030 (0.005)	0.029 (0.003)	0.029 (0.004)
GABA %CRLB	48.2 (25.8)	48.5 (34.6)	57.8 (36.1)	47.8 (23.2)
Glu %CRLB	4.7 (0.6)	4.8 (0.7)	4.9 (0.8)	4.9 (0.6)

Table 5.2: MRS characteristics. Group mean (standard deviation in parenthesis) voxel composition, spectrum signal-to-noise ratio (SNR), full width at half maximum (FWHM), and GABA and glutamate (Glu) relative Cramér-Rao lower bound (%CRLB) are shown for each spectroscopy measurement (before a-tDCS, after a-tDCS, before s-tDCS, after s-tDCS).

Relative Cramér-Rao Lower Bounds (CRLB) is commonly used as a quality-filtering criterion to identify and discard 'bad quality' data (Provencher 1993; Provencher 2001; Scholz *et al.* 2009; Kim *et al.* 2014; Antonenko *et al.* 2017). However, as recently highlighted in the work of Kreis (2016), this method induces a selection bias towards high concentration estimates and therefore leads to an over-estimation of the group mean metabolite concentration. This is particularly problematic for low concentration metabolites such as GABA. In order to avoid this methodological confound in this experiment, we used the alternative quality filtering method discussed in chapter 3. First, a high relative CRLB threshold of 100% was applied in order to exclude exceptionally unreliable measurements. Next, we calculated the expected relative CRLB value for all MRS measurements under

the assumption of a constant noise level across all measurements. We then excluded measurements for which the observed relative CRLB was exceptionally higher than the one predicted under this assumption. We used Pearson residuals to quantify the distance between the observed and predicted relative CRLB (threshold = 2). This quality filtering method allowed us to identify poor quality measurements irrespective to their concentration estimate and therefore avoid the selection bias of the standard method based on relative CRLB. Out of the 24 participants, 16 passed this quality check successfully for GABA, glutamate and creatine across all four MRS measurements. Twenty-three participants had at least one reliable baseline MRS measurement for all these metabolites (Table 5.1).




Figure 5.2: Magnetic Resonance Spectroscopy data. **A.** Example raw spectroscopy spectrum and LC model fit from one participant. The fitted LCModel (red) is overlaid on the raw data (black). The difference between the data and model (residuals) is shown at the top, and the baseline is shown at the bottom. The arrows represent the constituent peaks of glutamate (blue) and GABA (green). **B.** MRS voxel group overlap map. The MRS was centred on the left motor knob to include parts of the pre- and post-central gyrus (Yousry *et al.* 1997). This resulted in consistent positioning of the voxel in the same location in MNI space between individuals (light colours indicate overlap in voxel position in all 24 participants) and scanning sessions (red: anodal tDCS scan; blue: sham tDCS scan). **C.** The total creatine (tCr) concentration estimate was stable across all four measurements. Total creatine’s absence of significant variation across measurements ($F(1, 19) = 1.46$, $p = 0.24$) and good test-retest reliability ($ICC = 0.86$, $p < 10^{-3}$) enabled us to use it for internal referencing.

FMRIB’s automated segmentation tool (Patenaude *et al.* 2011), part of the FSL

software library (Smith *et al.* 2004), was used on the T1-weighted images to calculate the percentage of grey matter, white matter and cerebrospinal fluid present in the volume of interest (noted [GM], [WM] and [CSF] respectively) (Table 5.2). GABA and glutamate concentration estimates were corrected for the proportion of grey matter within the VOI (i.e. divided by $[GM]/([GM] + [WM] + [CSF])$), and creatine and phosphocreatine concentration estimates were corrected for the proportion of total brain matter within the VOI (i.e. divided by $([GM] + [WM])/([GM] + [WM] + [CSF])$).

Reliability of total creatine (creatine + phosphocreatine, noted tCr) concentration estimate across all four measurements was assessed using intra-class correlations (two-way random effects model with absolute agreement) implemented in SPSS (SPSS inc., version 24.0). An excellent degree of test-retest reliability was observed for tCr ($ICC = 0.86$, $p < 10^{-3}$), enabling us to use it for internal referencing. In all analyses, GABA and glutamate levels were therefore reported as ratios to tCr. Across the two baseline measurements, test-retest reliability was good for Glu:tCr ($ICC = 0.75$, $p = 0.001$) but lower for GABA:tCr ($ICC = 0.28$, $p = 0.26$). Thus, when evaluating the relationship between baseline metabolite concentration and computational parameters, we averaged the baseline concentration estimates across the two scanning sessions to obtain a more reliable estimate.

5.2.10.2 Diffusion-weighted images

Analysis of DW images was performed using tools from the FMRIB Software Library v5.0 (FSL; <https://fsl.fmrib.ox.ac.uk>) (Smith *et al.* 2004; Woolrich *et al.* 2009; Jenkinson *et al.* 2012). Preprocessing included the correction for susceptibility distortions, eddy current distortions and subject movement artefacts using TOPUP (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/topup>) (Andersson *et al.*

2003) and EDDY (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/eddy>) (Andersson & Sotiropoulos 2016). A diffusion tensor model was then fitted to the corrected DWI data using FDT (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FDT>) (Behrens *et al.* 2003; Behrens *et al.* 2007a) in order to generate individual Fractional Anisotropy (FA) and Mean Diffusivity (MD) maps.

Voxelwise statistical analysis of the FA data was carried out using TBSS (Tract-Based Spatial Statistics, <https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/TBSS>) (Smith *et al.* 2006b). In brief, all subjects' FA and MD images were first brain-extracted using BET (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/BET>) (Smith 2002) and then aligned into a common space using the nonlinear registration tool FNIRT (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FNIRT>) (Andersson *et al.* 2007). Next, the mean FA image was created and thinned to create a mean FA skeleton, which represents the centres of all tracts common to the group. Each subject's aligned FA data was then perpendicularly projected onto this skeleton for statistical analysis. The same procedure was applied to MD data. Visual observation of the raw DWI data and skeletonised FA and MD maps revealed unreliable measurements in the temporal lobes due to large signal dropouts. Therefore, voxels from the temporal lobes were excluded from further statistical analysis.

The significance of the association between individual differences in weight of the fast and slow systems and white matter microstructure (FA and MD) was tested using Permutation Analysis of Linear Models (PALM) (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/PALM>; version alpha-109) (Winkler *et al.* 2014; Winkler *et al.* 2016a). The advantage of this technique is that it allows carrying out voxelwise joint inference on FA and MD using a full nonparametric permutation testing and Fisher combination to calculate Family-Wise Error (FWE)-corrected significance level. Re-

sults were considered significant for $p < 0.05$ after FWE 2D threshold-free cluster extent correction and correction across contrasts. The location of significant clusters was determined using the JHU DTI-based white matter atlases (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases>) (Mori *et al.* 2005).

5.2.10.3 Resting state functional MRI

Analysis of resting state fMRI data was carried out using an independent component analysis (ICA) approach as implemented in MELODIC (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/MELODIC>) (Beckmann *et al.* 2005) and tools from FSL (Smith *et al.* 2004). The first 5 volumes of each fMRI scan were discarded to allow for T1 equilibration effects. Standard pre-processing steps were applied, which included motion correction (MCFLIRT) (Jenkinson *et al.* 2002), brain extraction (BET) (Jenkinson *et al.* 2005), B0 unwarping (FUGUE), spatial smoothing using a Gaussian kernel of full-width at half- maximum (FWHM) of 5 mm, and high-pass temporal filtering equivalent to 100 s. fMRI volumes were registered to the individual’s structural scan using boundary-based registration (BBR) (Greve & Fischl 2009) and then to standard space images using FNIRT (Andersson *et al.* 2007). Individual fMRI data were denoised using FIX (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FIX>, version v1.065 beta) trained on a subset of 24 scans taken from our dataset (Griffanti *et al.* 2014).

Pre-processed functional data containing 280 volumes for each time point (four time points per subject: pre/post, real/sham a-tDCS) and each subject ($n = 22$; [Table 5.1](#)) were temporally concatenated across subjects to create a single 4D data set that was decomposed into 25 components using ICA (Stagg *et al.* 2014; Bachtiar *et al.* 2015). All components were manually screened and non-physiological components were discarded. Informed decision was based on thresholded spatial maps ($Z > 4$, regional loci outside reasonable areas), as well as Fourier frequency decomposition of




Figure 5.3: Thirteen resting-state networks extracted from the 25-component analysis. All ICA spatial maps were converted to z-statistic images via a normalised mixture-model fit and thresholded at $Z > 4$. The resting state network highlighted in a red square (RSN 13₂₅) is the “hand” sensorimotor network and lies underneath the anode electrode.

the components’ time courses (shift toward high frequencies) (Beckmann *et al.* 2005; Griffanti *et al.* 2014). In total, 13 components were kept (Figure 5.3). They included three “visual” networks (maps 1₂₅, 2₂₅ and 3₂₅), the default-mode network (map 4₂₅), an “auditory” network (map 5₂₅), a “cerebellar” network (map 6₂₅), an “insular” network (map 7₂₅), an “executive control” network (map 8₂₅), two “frontoparietal” networks (maps map 9₂₅ and 10₂₅) and three “sensorimotor” networks (maps map 11₂₅, 12₂₅ and 13₂₅). The three sensorimotor networks reflected the somatotopic organisation of the primary sensorimotor cortices and superior cerebellum, with a predominantly leg (medial), face (inferior) and hand (lateral) component (maps 11₂₅, 12₂₅ and 13₂₅ respectively; Figure 5.3) (Van Den Heuvel *et al.* 2010).

Next, we estimated individual component spatial maps and corresponding time series for each subject and time point using dual regression (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/DualRegression>) (Beckmann *et al.* 2009a; Filippini *et al.* 2009). Network modelling was then carried out using the FSLnets toolbox (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FSLNets>) (Smith *et al.* 2015). Netmats were estimated using partial temporal correlations between components' time series, with a small amount of L2 regularisation ($\rho = 0.1$ in the Ridge regression netmats option in FSLnets, i.e. netmats5). Pearson correlation coefficients were converted into Z-statistics with Fisher's transformation before statistical analysis. The relationship between our computational parameters of interest (weights of the fast and slow systems) and the degree of functional connectivity between RSNs was investigated within a GLM framework. In order to minimise multiple comparisons, only edges involving the "hand" sensorimotor component (map 13₂₅) were included in the analysis. Non parametric testing was used using PALM (Winkler *et al.* 2014; Winkler *et al.* 2016a) with correction across contrasts.

5.3 Results

5.3.1 Computational model of sensorimotor memory formation

To formally test the hypothesis that reduced M1 inhibition biases temporal credit assignment towards longer timescales, right-handed healthy volunteers ($n = 24$, mean age = 25 years, SD = 3.50) adapted to 10° right-shifting prisms, during or after real or sham left M1 anodal tDCS (four sessions per individual, Figure 5.1). The PA paradigm was as before (O'Shea *et al.* 2017). Participants adapted gradually to prisms over 20 minutes (7 blocks, 120 trials, Figure 5.4B: E1-7). Leftward AEs were

probed in interleaved blocks during PA (7 blocks, 15 trials each, [Figure 5.4B](#): AE1-7). AE retention was assessed after a 10-minute rest delay (45 trials; [Figure 5.4B](#): AE8-10).

To extract the computation of interest (i.e. the estimate of the timescale of the visual disturbance), we built a variant of an influential state-space model of adaptation that posits two adaptive processes learning from performance error – one 'fast' that learns and forgets quickly and one 'slow' that learns and forgets slowly (Smith *et al.* 2006a). The net adaptation (during error-dependent learning and AE measurements) is the summed output of the two systems. During PA, error-dependent learning behaviour is relatively invariant across participants ([Figure 5.4B](#): E1-7), however individuals differ in how the consequent AE develops and decays ([Figure 5.4B](#): AE1-7). The original model cannot account for this variability because the dynamics of AE behaviour are dependent fully upon the dynamics of error correction. To remedy this, we modelled AE as a weighted sum of the two systems ([Figure 5.4](#)).




Figure 5.4: Computational modelling of motor behaviour during prism adaptation. **A.** Schematic representation of the weighted two-state model. **B.** Fit of the group mean prism adaptation data. The x-axis represents trial number and the y-axis pointing accuracy (rightward errors are positive). Black wedges indicate blocks throughout which prisms were worn. Group mean data (four sessions per participants, i.e. 96 datasets) is plotted in black (error bar = SEM). The fitted weighted two-state model is represented in red and the time courses of the fast and slow systems in green and blue respectively. **C.** Individual fits in the absence of stimulation. Each bar plot represents the relative weight (R_w) of one participant (estimated on the sham-tDCS conditions only), ranked in ascending order from left to right. Participant numbers are indicated on top of each bar (bold = Val/Met carriers, regular = Val/Val carriers, parenthesis = not genotyped). Individual model fits are plotted for the two extreme participants on this spectrum (lowest R_w is S7; highest R_w = S8). The individual with lowest R_w shows temporal dynamics of AE that more strongly resemble the fast (green) than the slow (blue) system, with steep slopes indicating a labile AE during adaptation, which declines in magnitude across adaptation to retention. In the individual with highest R_w , AE1-7 more strongly reflects the dynamics of the slow system (blue) than the fast (green), with flat slopes indicating stable AE during adaptation, which increases in magnitude across Adaptation to Retention.

Figure 5.4B shows this model provided a good fit to the group mean data ($R^2 = 99.68\%$), outperforming the original model after correcting for the benefit of two extra degrees of freedom ($\Delta AIC = -273.02$; full details in section 5.2.6). When fitting individuals' data, the learning and retention rates of both systems were kept constant for all. This approach reduced the number of free parameters, enabling individuals' AE to be modelled solely in terms of the two weights, one per system. Within this modelling framework, temporal credit assignment can be conceptualised as the relative weight assigned to the slow versus fast system, i.e. short versus long memory timescales. It is computed as: $R_w = w_s / (w_f + w_s)$. Theoretically, higher R_w predicts more stable AE.

5.3.2 Functional and structural connectivity correlates of memory persistence

The model quantitatively captures inter-individual variation in PA profile, based on a principled theory of how (hidden) computational states contribute to (overt) motor behaviour. This suggests its parameters may be sensitive to the underlying neural structures implementing these computations. To test this, we regressed the model parameters (w_f , w_s) against whole brain imaging measures of white matter microstructure and resting state functional connectivity. Figure 5.5 shows the structural and functional correlates of w_s across participants (w_f regressed out). Inter-individual variation in w_s correlated with the microstructural integrity of the left superior longitudinal fasciculus (predominantly branches II and III), a white matter pathway interconnecting frontoparietal cortex (Figure 5.5A; $p < 0.05$, FWE-corrected with 2D-cluster enhancement) (Thiebaut De Schotten *et al.* 2011). Across individuals, the higher the weight assigned to the slow system on AE, the stronger the coherence of white matter connectivity within this tract (i.e. higher fractional anisotropy and lower mean diffusivity).




A) Structural connectivity**B) Functional connectivity****C) 3D rendering of the structural and functional connectivity results**

Figure 5.5: Structural and functional connectivity correlates of memory persistence. **A.** White matter whose microstructure covariates with the weight of the slow system (w_s). Light blue shows the white matter FA skeleton in which statistical analysis (TBSS) was carried out. Dark blue indicates regions where FA is positively associated with w_s and MD is negatively associated with w_s (Fisher non-parametric combination, $p < 0.05$, FWE-corrected with 2D-cluster enhancement, w_f regressed out). Linear relationships are plotted underneath for illustration purpose. **B.** In the absence of stimulation, the degree of functional coupling between these two resting state networks (left: "hand" sensorimotor network, $Z > 4$, $x = -38$, $y = -23$, $z = 54$; right: right frontoparietal network, $Z > 4$, $x = 39$, $y = -54$, $z = 45$) was positively associated with the weight of the slow system ($p < 0.05$, FWE-corrected, w_f regressed out). The relationship is plotted underneath for illustration purpose. **C.** 3D rendering of the structural and functional connectivity correlates of the weight of the slow system. On this glass brain are overlaid: 1) the aforementioned white matter voxels (in blue) whose microstructure is associated with w_s , and 2) the "hand" sensorimotor network (in red, $Z > 4$) and right frontoparietal network (in green, $Z > 4$) whose functional connectivity is associated with w_s .

Analyses of resting state data asked whether, across individuals, there was any functional network for which the strength of coupling with the “hand” sensorimotor network (RSN 13₂₅; [Figure 5.3](#)) co-varied quantitatively with w_s (while controlling for w_f). [Figure 5.5B](#) shows that only one network out of 12 showed this pattern: the right frontoparietal attention network (RSN 9₂₅; [Figure 5.3](#)). Individuals with higher weight assigned to the slow system on AE measures during PA had greater functional coupling between the sensorimotor network and the right frontoparietal attention network in the resting state (edge 13₂₅-9₂₅; $p = 0.038$, FWE-corrected). These structural and functional relationships with w_s were independent, and remained significant when analyses tested for one while controlling for the other (both $p < 0.03$). This suggests that trait variation in resting state structural and functional connectivity of parietal-sensorimotor-premotor circuits influences the (long) time constant of sensorimotor memory and decay.

5.3.3 Sensorimotor cortex excitation:inhibition ratio causally determines memory persistence

Next, we tested our key mechanistic hypothesis: that excitation:inhibition balance in sensorimotor cortex causally determines whether AE persists or decays by biasing temporal credit assignment towards longer timescales. [Figure 5.6A](#) shows our neurocomputational hypotheses and one-way directional predictions about the causal relationships between sensorimotor cortex neurochemistry, temporal credit assignment, and behaviour at baseline (i.e. in the absence of excitatory stimulation). All the statistical analyses reported in this section were one-tailed given the *a priori* directional hypotheses set out in [Figure 5.6A](#).

First, we tested for a baseline relationship between sensorimotor cortex E:I and temporal credit assignment. Excitation:inhibition ratio was quantified across indi-

viduals as the relative concentration of glutamate to GABA in the left sensorimotor cortex, using high-field (7T) resting state MRS. Across individuals, we predicted a positive correlation between E:I and the relative weight parameter (R_w). Robust linear regression confirmed this (Glu:GABA $\times R_w$: $\beta = 0.42$, $p = 0.01$; Figure 5.6B). When evaluating the individual contribution of each neurotransmitter to this relationship (analyses controlled for the other metabolite), greater GABA levels were found to be associated with lower R_w ($\beta = -0.68$, $p < 10^{-4}$). By contrast, glutamate showed no association with R_w ($\beta = 0.05$, $p = 0.75$). In turn, as expected, higher R_w during PA predicted greater and more persistent AE at retention (AE magnitude: $\beta = 0.46$, $p = 0.01$; AE persistence: $\beta = 0.76$, $p < 10^{-3}$; Figure 5.6B).




Figure 5.6: Predictions and correlations of MRS, computational and retention measures at baseline. A. Based on previous findings (Galea *et al.* 2010; O'Shea *et al.* 2017), we predicted greater M1 E:I ratio to be associated with greater relative contribution of the slow system to the AE during adaptation, which, in return, should predict greater and more persistent AEs at retention.

(Legend continued on next page)

B. Analysis of baseline (i.e. in the absence of stimulation) inter-individual co-variation in these three levels confirmed the predictions. Across individuals, higher Glu:GABA was associated with greater R_w ($\beta = 0.42$, $p = 0.01$). Individual contribution of GABA and Glu to this relationship is plotted on the side. In return, greater R_w during PA was associated with greater ($\beta = 0.46$, $p = 0.01$) and more persistent ($\beta = 0.76$, $p < 10^{-3}$) AEs after a 10 minutes break. Colour code represents the contribution of individual data points to the linear regression.

To test the hypothesis that R_w depends *causally* on E:I, we intervened within sensorimotor cortex with excitatory stimulation (a-tDCS) during PA. This has previously been shown to reduce GABA concentration (Stagg *et al.* 2009; Bachtiar *et al.* 2015; Antonenko *et al.* 2017). If E:I causally determines R_w , then stimulation during PA should increase R_w and consequently increase AE at retention (Figure 5.7A). If the computation of R_w is critical for this effect to arise, the same should not be observed when stimulation is delivered before PA (offline tDCS condition). In other words, across individuals, the magnitude of stimulation-induced neurochemical change (Δ E:I) should predict the computational change (ΔR_w) induced by online (but not offline) a-tDCS, which should, in return, predict the behavioural change at retention (ΔAE).

At the group mean level, stimulation did not change GABA concentration ($p = 0.31$; Figure 5.10), but stimulation during PA did increase AE at retention ($p = 0.02$, Cohen's $d = 0.4640$; Figure 5.9A), while stimulation prior to PA did not ($p = 0.64$; Figure 5.9A), all three findings replicating previous work (full details in section 5.3.5; O'Shea *et al.* 2017). Critically, analyses of inter-individual co-variation in the effect of stimulation at the neurochemical, computational and behavioural level confirmed the predictions (Figure 5.7A). Across individuals, the greater the induced increase in Glu:GABA, the greater the increase in R_w when a-tDCS was applied during PA ($\beta = 0.57$, $p = 0.008$, controlling for Δ Glu:GABA occurring with sham tDCS; Figure 5.7C). As expected, the stimulation-induced change in computation predicted

behavioural change at retention. The greater the induced increase in R_w , the greater the induced increase in AE at Retention (Δ AE magnitude: $\beta = 0.41, p = 0.02$; Δ AE persistence: $\beta = 0.33, p = 0.05$; Figure 5.7).




Figure 5.7: Predictions and correlations of stimulation-induced change in MRS, computational and retention measures. **A.** If the relationship described in Figure 5.6 are causal, then across individuals, the magnitude of stimulation-induced increase in E:I ratio should relate to how much R_w increase, which should predict subsequent increase in AE at retention. **B.** Regression analyses confirmed these predictions (Δ GABA:Glu $\times R_w$: $\beta = 0.57, p = 0.008$; Δ AE magnitude: $\beta = 0.41, p = 0.02$; Δ AE persistence: $\beta = 0.33, p = 0.05$). Colour code represents the contribution of individual data points to the linear regression.

The analysis of the individual relationships of Δ GABA and Δ Glu with ΔR_w (analysis controlling for the a-tDCS-induced change in the other metabolite and sham tDCS-induced change in Glu:GABA) revealed the same pattern found at baseline. Stimulation-induced decrease in GABAergic tone was associated with increase in R_w ($\beta = -0.53, p = 0.03$) but changes in glutamatergic tone showed no such association ($\beta = 0.28, p = 0.32$). When stimulation was applied prior to PA, the a-tDCS-induced

$\Delta E:I$ did not correlate with ΔR_w ($\beta = -0.20; p = 0.43$). That induced $\Delta E:I$ increases AE retention, only when applied concurrent with (but not prior to) PA, implies causal dependence of stimulation on this computational state (R_w). This confirms our mechanistic prediction: that higher sensorimotor cortex excitation biases temporal credit assignment towards longer timescales during PA, causing lasting memory.

5.3.4 Enhanced memory persistence via sensorimotor cortex stimulation is modulated by genotype

Individuals with a common polymorphism in the gene coding for brain derived neurotrophic factor (BDNF val66met) exhibit reduced behavioural and neural markers of motor cortical plasticity (Kleim *et al.* 2006; McHughen *et al.* 2011; Joundi *et al.* 2012). Plastic enhancement of motor skill learning via a-tDCS is also reduced in Met allele carriers (Fritsch *et al.* 2010). The polymorphism causes a partial reduction in activity-dependent BDNF secretion (Egan *et al.* 2003; Chen *et al.* 2005), a factor involved in long-term potentiation (Minichiello 2009). Augmentation of BDNF-dependent synaptic plasticity is a candidate mechanism of action of sensorimotor cortex a-tDCS in mice and humans (Fritsch *et al.* 2010). Here we tested the prediction that this genotype would modulate the cognitive state dependent effect of stimulation on AE retention. That is, since the memory enhancement effect of sensorimotor cortex a-tDCS is activity-dependent (requires concurrent computation of R_w), the impact of stimulation applied during PA should vary with genotype, while stimulation at rest prior to PA may not. Genotyping was acquired for 21/24 participants after all other data were analysed.




Figure 5.8: The BDNF Val66Met polymorphism influences the effect of anodal tDCS on retention. **A.** The change in prism AE magnitude at retention between anodal and sham tDCS (a value > 0 means an increased in AE) is plotted as a function of experiment (online/offline) and BDNF polymorphism (Val/Val in grey, Val/Met in white). Each dot represents an individual and the error bars represent the SEM. BDNF polymorphism influenced the a-tDCS effect only when stimulation was delivered during but not before PA. Val/Val carriers ($n = 15$) but not Val/Met carriers showed the expected increase in AE with online a-tDCS. Asterisk indicates significant effects (($*$): $p = 0.053$; ($**$): $p < 0.01$). **B-C.** Val/Val carriers but not Val/Met carriers responded to online M1 a-tDCS. The x-axis represents trial number and the y-axis pointing accuracy in reference to baseline accuracy (i.e. change from baseline, positive values indicate rightward errors). Pointing accuracy in Val/Val carriers (**B**, $n = 15$) and Val/Met carriers (**C**, $n = 6$) is plotted when anodal (red) or sham (blue) was applied to the left M1 during prism adaptation (online experiment) (group mean ± 1 SEM). Black wedges indicate blocks throughout which prisms were worn (CLP trials).

In agreement with the known allele distribution in the Caucasian population (Verhagen *et al.* 2010), 6/21 participants (28.6% of our sample) carried the Met allele. Mean AE at retention was analysed across all four PA sessions (anodal/sham tDCS $\times$ before/during PA). The interaction of genotype with cognitive state and stimulation was marginally significant ($F(1,19) = 4.3$, $p = 0.053$, 2-tail, $\eta_p^2 = 0.18$;

Figure 5.8). Follow-up ANOVAs suggested that BDNF polymorphism influenced the behavioural response to stimulation when applied during PA (stimulation $\times$ genotype: $F(1, 19) = 9.24$, $p < 0.01$, 2-tail, $\eta_p^2 = 0.327$) but not when applied prior to PA ($F(1, 19) = 0.04$, $p = 0.84$, 2-tail). As predicted, stimulation during PA significantly enhanced AE retention in Val-Val carriers (anodal-sham: $t_{(14)} = -3.51$, $p < 0.01$) but not in Val-Met carriers ($t_{(5)} = 1.30$, $p = 0.25$). These findings implicate activity-dependent BDNF secretion in stimulation-induced enhancement of AE persistence. This supports the hypothesis that augmentation of BDNF-dependent synaptic plasticity is a contributory mechanism mediating behavioural plasticity induction via sensorimotor cortex anodal tDCS. Prior evidence was confined to motor skill learning. Our results extend this to sensorimotor adaptation.

5.3.5 Supplementary analyses (group mean effects)

5.3.5.1 Statistical analysis of motor behaviour during and after prism adaptation


The online and offline experiments aimed to replicate the finding that left M1 a-tDCS would enhance AE consolidation when applied during but not before PA. A group of 24 participants underwent PA and tDCS in a repeated measures design, in four separate sessions, at least 1 week apart, in which they received anodal or sham tDCS either before or during PA. RM ANOVA considered the within-subjects of tDCS (anodal versus sham) and block separately for each experimental condition (online and offline).

- **Prism exposure (Blocks E1-7)**

Accuracy on CLP improved progressively as adaptation developed in the online (Blocks E1-7: $F(3.24, 74.44) = 132.24$, $p < 10^{-3}$; Figure 5.9A) and offline

(Blocks E1-7: $F(1.61, 2.21) = 89.51$, $p < 10^{-3}$; Figure 5.9B) experiments. No other effects were significant, indicating that neither online nor offline a-tDCS significantly influenced block mean pointing accuracy on CLP.

A) Left M1 a-tDCS during PA enhances AEs at retention



B) Left M1 tDCS before PA has no effect on behaviour




Figure 5.9: Behavioural effect of left M1 a-tDCS. The x-axis represents trial number and the y-axis pointing accuracy in reference to baseline accuracy (i.e. change from baseline). Positive values represent a rightward shift and negative values leftward shift. Pointing accuracy in healthy volunteers ($n = 24$) is plotted (group mean ± 1 SEM) when anodal (red) or sham (blue) was applied to the left M1 during PA (A) or before PA (B). Black wedges indicate blocks throughout which prisms were worn (CLP trials). During prism exposure (E1-7), participants saw the outcome of each trial, so they could gradually correct their errors. The AE was measured without visual feedback (AE1-7, shaded light grey). Similar to previous report (O'Shea *et al.* 2017), relative to sham, anodal tDCS increased prism AE at retention when applied during ($p = 0.02$, $Cohen'sd = 0.46$) but not before ($p = 0.64$). Asterisk indicates significant difference between the anodal and sham online condition ($p < 0.05$).

- **Adaptation phase (Blocks AE1-7)** Interleaved with exposure blocks, prism AE was measured as adaptation developed. Prism AEs stabilised progressively

over time in both the online (Blocks AE1-7: $F(3.95, 90.78) = 11.07$, $p < 10^{-3}$; [Figure 5.9A](#)) and offline (Blocks AE1-7: $F(3.324, 76.456) = 15.408$, $p < 10^{-3}$; [Figure 5.9B](#)) experiments. However, no other effects were significant, indicating that neither online nor offline a-tDCS significantly influenced motor behaviour on OLP during the adaptation.

- **Retention phase (Blocks AE8-10)** Following PA and a 10 min rest period, AE was assessed. For the online experiment, a RM ANOVA with within-subjects factors of tDCS (anodal vs. sham) and blocks (AE8-10) confirmed the prediction of enhanced AE magnitude (main effect of tDCS: $F(1, 23) = 4.38$, $p = 0.02$, 1-tail, mean difference = -1.10 degrees, SEM = 0.53, Cohen's $d = 0.46$; [Figure 5.9A](#)). As expected, this effect was not found in the offline experiment ($F(1, 23) = 0.22$, $p = 0.64$, 2-tail, mean difference = -0.294 degrees, SEM = 0.6221; [Figure 5.9B](#)).

5.3.5.2 Statistical analysis of the neurochemical effect of left M1 a-tDCS

RM ANOVAs considered the within-subject factors of time (pre vs. post tDCS) and tDCS (anodal vs. sham) separately for each metabolite to evaluate the neurochemical effect of the stimulation. At the group level, a-tDCS did not affect GABA ($F(1, 15) = 1.10$, $p = 0.31$), glutamate ($F(1, 19) = 2.31$, $p = 0.15$) or Glu:GABA ($F(1, 15) = 0.25$, $p = 0.62$) levels, as illustrated by no interaction 'time $\times$ tDCS' ([Figure 5.10](#)). In other words, at the group level, we found no evidence for the previously reported reduction in GABAergic inhibition with a-tDCS (Stagg *et al.* 2009; Bachtiar *et al.* 2015; Antonenko *et al.* 2017). Nevertheless, the neurochemical response to a-tDCS showed large inter-individual variability, which enabled us to address our question of interest: whether individuals with larger a-tDCS-induced increase in E:I ratio would also show larger increase in R_w when a-tDCS is applied during PA (results reported above).




Figure 5.10: No group mean effect of stimulation on metabolite concentrations. Bar graphs represent the GABA:tCr ratio (A) and Glu:tCr ratio (B) within the MRS voxel before and after real (in red) or sham (in blue) left M1 a-tDCS inside the scanner. Individual participants are represented as lines overlaid on the bar graphs. Only individuals with complete dataset (i.e. satisfactory data quality for all four time points; see Table 5.1) are represented. RM ANOVAs revealed no significant interaction 'time $\times$ tDCS' for either metabolite (both $p > 0.15$). Note, however, the large inter-individual variability in neurochemical response to a-tDCS, exploited in Figure 5.7B.

5.3.5.3 Statistical analysis of the a-tDCS effect on resting-state networks

Informed by the baseline relationship between w_s and functional connectivity between the “hand” sensorimotor RSN and right frontoparietal RSN (edge 13₂₅-9₂₅), we asked whether the strength of this edge was influenced by M1 a-tDCS in the scanner. Repeated measure ANOVA revealed a main effect of scanning session ($F(1, 21) = 5.70$, $p = 0.03$) but no main effect of time ($F(1, 21) = 0.10$, $p = 0.76$) or interaction 'time $\times$ stimulation' ($F(1, 21) = 0.55$, $p = 0.47$). Additionally, inter-individual variations in the a-tDCS change in the edge 13₂₅-9₂₅ did not co-vary with inter-individual changes in w_s ($\beta = 0.08$, $p = 0.74$, while controlling for the fast system). These null-results provide further support for the cognitive state dependency of the a-tDCS effect. The effect of a-tDCS on the functional coupling between the two networks is likely to require engaging individuals in the PA task, which was not the case inside the scanner.

A previous study reported increased functional connectivity within M1 with M1 a-tDCS (Bachtiar *et al.* 2015). In an attempt to replicate this result, we centred a $9 \times 9 \times 9 \text{ mm}^3$ region of interest (ROI) on the peak coordinates of the “hand” sensorimotor RSN within the left M1 ($x = -37, y = -24, z = 62; Z = 9.4$) and extracted the mean network strength for each subject, stimulation condition, and time point. Unlike in Bachtiar *et al.* (2015)’s study, repeated measures ANOVA showed no significant main effect of stimulation ($F(1, 21) = 0.39, p = 0.54$) or time ($F(1, 21) = 1.10, p = 0.31$) and no significant stimulation $\times$ time interaction ($F(1, 21) = 0.04, p = 0.84$).

5.4 Discussion

Adaptation is a fundamental property of the nervous system that underwrites the maintenance of successful actions across the lifespan (Shadmehr *et al.* 2010). Computational theories posit that individuals adapt to perturbations (e.g. lateral displacement of the visual field by prisms) by forming an internal model of it, used to adjust motor output and restore accuracy (Smith *et al.* 2006a). Once acquired, what determines whether adaptation memory persists or decays? The present study investigated this question from a computational and neurochemical perspective. We recently showed (and replicated in the present study) that excitatory stimulation (a-tDCS) of sensorimotor cortex during but not before PA increases retention of the AE (O’Shea *et al.* 2017). This timing-dependent effect indicates that understanding the information processing implemented within sensorimotor cortex and interacting with the stimulation is key to understanding the mechanism by which sensorimotor memories are formed and enhanced. Here, we hypothesised that sensorimotor cortex excitation:inhibition during PA might reflect the solution to the temporal credit assignment problem, such that lower inhibition would bias estimates towards longer

timescales, and therefore longer-lasting AEs. To formally test this hypothesis, we inferred hidden computational states from behaviour using a variant of an influential state-space model (Smith *et al.* 2006a) that allowed us to quantify the relative contribution of short versus long memory timescale to the internal representation of the perturbation (chapter 4). This weighted two-state model outperformed the original model and provided a good fit to group mean and individual data. Across individuals, greater E:I ratio (quantified as Glutamate:GABA concentration using 7T MRS) was associated with greater contribution of the slow system during PA, which in return predicted greater and more persistent AEs at retention. To establish causality, the effect of excitatory a-tDCS was assessed on all three levels (physiology, computation, behaviour). This resulted in correlated increases in E:I, relative contribution of the slow system, and retention.

MR spectroscopy offers a way to measure metabolite concentrations in the living brain. The non-invasiveness of this technique comes at the cost of a relatively low signal-to-noise (SNR) ratio. In this study, good SNR was obtained by using ultra high-field (7T) MRS and by acquiring data from a large $2 \times 2 \times 2$ cm³ MRS voxel centred on the region of interest. Due to the size of the voxel, adjacent regions of somatosensory cortex (S1) were also included in the measure of M1 metabolites. This is a common methodological limitation of MRS studies (Stagg *et al.* 2009; Lunghi *et al.* 2015; Antonenko *et al.* 2017; Ip *et al.* 2017; Kolasinski *et al.* 2017; O’Shea *et al.* 2017). Although we cannot rule out the contribution of S1 to our results, M1 is likely to play a predominant role because of a convergence of studies implicating it in the consolidation of adaptation memory (Wise *et al.* 1998; Li *et al.* 2001; Paz *et al.* 2003; Paz *et al.* 2005; Hadipour-Niktarash *et al.* 2007; Landi *et al.* 2011; Leow *et al.* 2016). It is not our intention however to suggest that only M1 is involved in the implementation of temporal credit assignment. As illustrated in the structural

and functional connectivity results, the weight assigned to the slow system appears to involve the interaction of M1 with at least the posterior parietal cortex and ventral premotor cortex, two regions known to be involved in PA (Kurata & Hoshi 1999; Clower & Boussaoud 2000; Pisella *et al.* 2004; Newport & Jackson 2006; Newport *et al.* 2006; Danckert *et al.* 2008; Luauté *et al.* 2009). A parsimonious interpretation of the results is therefore that the level of inhibition within M1 is a causal factor in the computation of the timescale over which to retain the solution of the adaptation, but that it may also involve other distant regions.

The relationships between M1 inhibition and R_w could result from two functionally distinct forms of GABAergic signalling: phasic or tonic. Phasic GABAergic activity refers to the release of presynaptic GABA into the synaptic cleft due to the depolarisation of inhibitory neurons (Mody 2001; Bachtiar & Stagg 2014). It is involved in the generation of rapid inhibitory postsynaptic currents (IPSCs) via the activation of GABA_B receptors (Maffei *et al.* 2017). In addition, extracellular GABA also has a longer-lasting neuromodulatory role via tonic signalling on extra-synaptic GABA_A receptors (Glykys & Mody 2007). This form of tonic inhibition affects neuronal excitability via an action on the tonic conductance, which is distinct from the time-dependent phasic inhibitory signalling (Mody 2001). It has an important regulatory influence on synaptic plasticity in that greater tonic inhibition is typically associated with less long-term potentiation (LTP) (Wigström & Gustafsson 1983; Chapman *et al.* 1998; Levkovitz *et al.* 1999; Bütefisch *et al.* 2000; Yoshiike *et al.* 2008; Martin *et al.* 2010). Extrasynaptic tonic GABA represents approximately 70% of the total GABA in the cortex (Petroff 2002) and, unlike vesicular phasic GABA, it is not bound to macromolecules, which makes it more easily detectable with MRS (Stagg *et al.* 2011c). Although these observations indicate that the resting MRS-GABA signal measured in this study reflects predominantly tonic extrasynaptic GABA_A inhibition

(Stagg *et al.* 2011b; Stagg *et al.* 2011c), it is currently nearly impossible to disentangle the contribution of the two neurotransmitter pools with the resolution of this technique. Studies considering TMS intra-cortical inhibitory measures as indirect proxy of GABA_A and GABA_B inhibition have reported modulatory effects of a-tDCS on both types of signalling (Hummel *et al.* 2005; Antal *et al.* 2010; Tremblay *et al.* 2013; Amadi *et al.* 2015). We therefore suggest that the baseline relationship between MRS-GABA and R_w is likely to rely on individual trait variation in GABA_A tone, while the relationship between Δ MRS-GABA and ΔR_w is likely to implicate state variation in both GABA_A and GABA_B. The latter might contribute to explaining why although M1 a-tDCS has been shown to reduce MRS-GABA measures beyond the period of the stimulation (for up to 30 minutes) (Bachtiar *et al.* 2015), it had no effect on behaviour when applied before PA (O'Shea *et al.* 2017). Another account for the timing-dependent interaction between a-tDCS and sensorimotor memory is that offline a-tDCS might involve metaplasticity (Stagg *et al.* 2011d), a set of mechanisms engaged to counteract the effect of excitatory stimulation in order to maintain neural activity within a normal range (Lang *et al.* 2004). These mechanisms might take place during PA in the offline tDCS condition and suppress the excitatory effect of a-tDCS.

Identification of individual predictors of responsiveness to stimulation is crucial for both the mechanistic understanding of the effect of tDCS and the tailoring of interventions on an individual basis. In the present study, the cognitive state dependency of the stimulation-induced memory enhancement was influenced by individuals' BDNF val66met polymorphism. Homozygote individuals (Val/Val) responded to online but not offline a-tDCS while heterozygotes individuals (Val/Met) responded to neither online nor offline a-tDCS. This result extends the known role of activity-dependent BDNF secretion in motor skill learning and its modulation by a-tDCS (Fritsch *et*

al. 2010) to another paradigm: sensorimotor adaptation. It is unclear however how BDNF polymorphism interacts with the neuro-computational mechanism reported above. For example, it might be the case that because of altered activity-dependent synaptic plasticity, Met carriers show a different relationship between neurochemical and computational changes (e.g. a greater reduction in GABA is needed to induce the same increase in R_w in Met carriers). The small number of Met carriers in our experimental sample ($n = 6$) prevented us from formally asking this question. Future research beyond the exploratory analysis reported here should therefore consider screening participants for their BDNF polymorphism at inclusion in order to obtain balanced sample sizes between the two groups (Val/Val and Val/Met) and formally compare the mechanisms involved.

5.5 Conclusions

The data presented in this chapter demonstrate a causal neuro-computational mechanism by which trait and state variation in sensorimotor GABAergic inhibition causally determines memory persistence/decay. It has important clinical implications for post-stroke rehabilitation of spatial neglect because it suggests that excitatory stimulation of the sensorimotor circuit during PA might boost the contribution of the slow system to prism AEs, which may result in increased persistence of the therapeutic effects.

Chapter 6

Sensorimotor memory formation in the ageing human brain

The previous chapter presented a causal mechanistic model of how human sensorimotor memories are formed, retained and enhanced by M1 anodal tDCS. In the present chapter, the same neuro-computational model is tested in the context of normal ageing.

6.1 Introduction

We demonstrated in the previous chapter ([chapter 5](#)) that the level of excitation:inhibition (E:I) in primary motor cortex (M1) causally determines the timescale of adaptation memory in young healthy adults. Specifically, lower M1 GABAergic inhibition was associated with greater contribution of long memory timescales to the prism after-effect (AE). Reducing inhibition via anodal transcranial direct current stimulation (a-tDCS) within M1 during prism adaptation (PA) enhanced retention of sensorimotor AE by changing a computational parameter reflecting temporal credit assignment (relative weight, R_w).

The primary clinical motivation for better understanding memory formation during PA is to maximise persistence of the AE in the context of post-stroke neglect rehabilitation. As outlined in the general introduction of this thesis ([chapter 1](#)), prism AE can transfer to improve cognitive deficits in visuospatial neglect after right hemisphere brain damage (Rossetti *et al.* 1998; Frassinetti *et al.* 2002; Serino *et al.* 2009; Barrett *et al.* 2012; Jacquin-Courtois *et al.* 2013). Although little is understood about the neural mechanisms underlying such cognitive improvements, the induction of a sensorimotor AE (i.e. leftward shift in pointing behaviour) appears to be an essential prerequisite for the therapeutic effect to arise after PA (Serino *et al.* 2006; Newport & Schenk 2012; Rode *et al.* 2015; O'Shea *et al.* 2017). In light of the previous findings ([chapter 5](#); O'Shea *et al.* 2017), M1 anodal tDCS appears to be a promising tool that could be used to maximise retention of the AE after single sessions of PA, which may lead to greater duration and/or efficacy of prism therapy for neglect rehabilitation (O'Shea *et al.* 2017).

An important point to consider however is the age difference between the experimental sample studied in the previous chapter ([chapter 5](#), mean age = 25) and the clinical population of interest (e.g. [chapter 8](#), mean age = 61). Stroke affects predominantly elderly adults (Feigin *et al.* 2014), with risk of having a stroke approximately doubling every decade after the age of 55 (Wolf *et al.* 1992; Brown *et al.* 1996). Because of functional re-organisations occurring with normal ageing (Gao *et al.* 2013; Heise *et al.* 2013; Wolpe *et al.* 2016; Porges *et al.* 2017a), findings in young adults may not be translatable directly to older adults or stroke patients. It is therefore crucial to investigate whether the relationship between M1 excitation:inhibition balance and AE persistence ([chapter 5](#)) is maintained in older adults.

Ageing has been associated with a general down-regulation of GABAergic inhibition (Gao *et al.* 2013; He *et al.* 2016; Porges *et al.* 2017a; Porges *et al.* 2017b) that is relatively well-documented in the primary motor cortex (Peinemann *et al.* 2001; Oliviero *et al.* 2006; Heise *et al.* 2013). A recent cross-sectional study linked this age-related reduction in tonic M1 inhibition to poorer dexterous motor performance on an alternating finger-tapping task and simple reaction time task (Heise *et al.* 2013). Although the causal status of this relationship has yet to be determined, the authors speculated that the behavioural impairments in older adults are likely to result from the neurochemical alterations accompanying ageing. If the neurochemical mechanism of AE persistence identified in young adults in the previous chapter ([chapter 5](#)) also applies in older adults, it makes the 'paradoxical' prediction that older age should be associated with greater AE retention, due to lower M1 GABA levels. Previous research investigating the effect of ageing on adaptation showed that although older adults demonstrate deficits during the exposure phase (i.e. slower error reduction) (Fernández-Ruiz *et al.* 2000; Buch *et al.* 2003; Bock 2005; Anguera *et al.* 2011; Huang & Ahmed 2014; Panouillères *et al.* 2015), the AEs following removal of the sensorimotor perturbation are similar to those of young adults (Roller *et al.* 2002; Buch *et al.* 2003; Bock 2005; Hegele & Heuer 2010; Panouillères *et al.* 2015) or even increased (Fernández-Ruiz *et al.* 2000; Nemanich & Earhart 2015). This pattern of behavioural performance fits well with our prediction.

Here, we asked whether the reduction in M1 GABAergic inhibition accompanying ageing explains changes in AE persistence. Across individuals, older age was found to be associated with greater long-term retention of the AE and lower M1 GABA. A mediation analysis suggested that the latter explained the former. This confirms, in a sample of older adults, the crucial role of M1 tonic inhibition in the neural mechanisms underlying AE persistence. Second, we applied left M1 anodal tDCS during

PA and tested its effect on AE persistence. To date, only 18 participants underwent this experiment but preliminary results indicate a beneficial effect on long-term retention, especially in individuals with low M1 excitation:inhibition. All together, these findings indicate preserved neurochemical mechanisms underlying AE persistence in older adults and suggest a good translatability of the findings of the previous and current chapters to the clinical population of interest.

6.2 Material and methods

6.2.1 Participants

In total, 27 right-handed male individuals aged between 49 and 81 years (mean age = 67.8 years, SD = 8.4) without any personal or family history of neurological or psychiatric disease participated in the experiments reported in this chapter. Similar to [chapter 5](#), the gender selection criterion was applied to avoid potential confounds related to variations in neurotransmitter concentration with the menstrual cycle in women (Smith *et al.* 1999; Epperson *et al.* 2002). There were three experiments:

- The first experiment (labelled 'Baseline PA' in [Table 6.1](#)) measured PA and 24-hour retention of the AE in all participants ($n = 27$). The aim was to investigate, across participants, age-related changes in PA performance and AE persistence. Adaptation data collected in these participants were also compared to the PA data of the young healthy adults tested in the previous chapter ($n = 24$, mean age = 25.03 years, SD = 3.50; [chapter 5](#)). For the purpose of this comparison, the averaged PA data over the two sham sessions were used for the young healthy adults.
- The second experiment (labelled 'MRS' in [Table 6.1](#)) was conducted in 22 of

these participants (mean age = 66.7 years, SD = 8.8) and involved magnetic resonance spectroscopy (MRS) measurements in a volume of interest centred on the left M1. In a sub-sample of 20 of these participants, MRS data were also collected from an anatomical control voxel centred on the bilateral visual cortex. The aim was to investigate, across participants, age-related changes in GABA and glutamate concentration in these 2 regions of interest, and to link these to inter-individuals differences in AE persistence.

Participant	Age	Baseline PA	MRS	tDCS + PA
1	81.1	✓	(✓)	✓
2	78.7	✓	(✓)	✓
3	80.0	✓	✓	
4	75.9	✓	✓	✓
5	74.6	✓	✓	✓
6	72.3	✓	✓	✓
7	68.7	✓	✓	✓
8	66.9	✓	✓	✓
9	64.6	✓	✓	✓
10	53.3	✓	✓	
11	68.9	✓	✓	✓
12	80.7	✓		
13	70.4	✓	✓	✓
14	66.0	✓	✓	
15	61.7	✓	✓	✓
16	68.3	✓	✓	
17	58.5	✓	✓	✓
18	53.4	✓	✓	
19	64.9	✓	✓	
20	61.7	✓	✓	✓
21	49.2	✓	✓	
22	61.6	✓	✓	✓
23	65.0	✓		✓
24	65.6	✓		✓
25	75.4	✓		✓
26	74.1	✓		✓
27	67.8	✓		✓

Table 6.1: Participation in each of the three experiments. Parenthesis indicate incomplete MRS data (i.e. M1 voxel only).

- The third experiment (labelled 'tDCS + PA' in [Table 6.1](#)) was conducted in 18 of these participants (mean age = 69.2 years, SD = 6.7). They underwent four behavioural sessions aiming to characterise the effect of left M1 anodal tDCS on PA in this population. Data collection for this experiment is still going because the target sample size is 24 in order to match the sample size of the experiment in young healthy adults reported in [chapter 5](#).

Details of the inclusion of the 27 healthy elderly participants in each of the three experiments (Baseline PA, MRS, PA + tDCS) is provided in [Table 6.1](#). Additionally MR spectroscopy data from 5 neglect patients (1 female, mean age = 60.9 years, SD = 10.4) tested in [chapter 7](#) & [8](#) are reported in this chapter. The aim was to compare GABA and glutamate levels in M1 between neglect patients and age-matched healthy controls. Recruitment information and characterisation of patients' baseline behavioural performance are detailed in [chapter 8](#). All participants gave written informed consent as approved by the Oxford A Research Ethics Committee (REC reference number: 12/SC/0556). No participant reported any side effect from the experimental procedure.

6.2.2 Experimental procedures

6.2.2.1 Baseline prism adaptation experiment

All healthy elderly participants ($n = 27$) first underwent 2 'baseline' behavioural sessions 24 hours apart from one another. PA was performed on day 1 and long-term retention was assessed on day 2. The PA protocol was as previously described ([Chapters 4](#) & [5](#)) and was run on the automated PA set-up I designed and built (see [4.2.3](#)).

There were 2 types of trials: closed-loop pointing (CLP) and open-loop pointing

(OLP). On CLP trials, participants made fast speeded reaching movements to point with their right index finger at left and right visual targets in counterbalanced order. Both accuracy and speed were emphasised. Visual feedback was limited to the last two thirds of the reaching movement in order to limit strategic adjustments and in-flight error correction (Redding & Wallace 1996; O’Shea *et al.* 2014; Inoue *et al.* 2015; O’Shea *et al.* 2017). On OLP trials, participants pointed at a comfortable speed to a visual target located in front of them. The target was only visible at the reach onset. Visual feedback of the reaching movement, terminal error and return movement were occluded by an opaque shutter, hence requiring participants to rely on proprioception to guide their movement.

Participants adapted gradually to 10° right-shifting prisms (20 minutes, 100 CLP trials over 6 blocks, E1-6) and leftward AEs were probed in interleaved blocks (15 OLP trials each, AE1-6). Retention of the AE was assessed after a 10-minute rest delay and after 24 hours (45 OLP trials in each time point).

6.2.2.2 MR spectroscopy experiment

MR spectroscopy data were acquired in 22 elderly healthy individuals and 5 neglect patients. Details of the full MRI scanning protocol is provided in [chapter 7](#). Recruitment information and characterisation of baseline behavioural performance of neglect patients are detailed in [chapter 8](#) ([Figure 7.1](#)). MRS data were acquired at 3T from two $2 \times 2 \times 2$ cm³ volumes of interest (VOI) in two consecutive acquisitions. This lower field strength (compared to the 7T scanner used in [chapter 5](#)) was used in order to lower the number of exclusion criteria, making it easier to recruit and test older adults and stroke patients. The first VOI (labelled M1) was centred on the left motor knob (Yousry *et al.* 1997) and included parts of the pre- and post- central gyrus ([Figure 6.1A](#), green voxel). The second (anatomical control) VOI (labelled V1) was

centred bilaterally on the calcarine sulcus in the occipital lobe (visual cortex) (Engel *et al.* 1997; Lunghi *et al.* 2015; Ip *et al.* 2017) in healthy elderly adults (Figure 6.1A, blue voxel). In neglect patients ($n = 5$), the visual voxel was shifted to the left hemisphere and centred on the left calcarine sulcus in order to avoid eventual lesioned neighbouring tissue in the right hemisphere.

B0 shimming was performed using a GRE-SHIM (64×4.2 mm axial slices, TR = 862.56 ms, TE1/2 = 4.80/9.60 ms, flip angle = 12 deg, FOV = 400, scan duration = 63 secs). MR spectroscopy data were acquired using the semi-LASER sequence (TR/TE = 4000/50 ms, 64 scan averages, scan time = 264 secs) (öz & Tkáč 2011; Bank *et al.* 2015).

6.2.2.3 Brain stimulation experiment

Eighteen older adults participated in a tDCS experiment involving two pairs of behavioural sessions (i.e. four behavioural sessions in total). The two pairs of behavioural sessions were separated by at least one week and their order was counter-balanced across participants. Each pair of behavioural sessions was similar to the baseline PA experiment. On day 1, participants adapted to 10° right-shifting prisms (20 minutes, 120 CLP trials over 7 blocks, E1-7) and leftward AEs were probed in interleaved blocks (15 OLP trials each, AE1-7). During the 20 minutes of PA, participants received either anodal or sham left M1 tDCS. Retention of the AE was assessed (45 OLP trials) after a 10-minutes rest delay (45 OLP trials) and after 24 hours.

6.2.3 Transcranial direct current stimulation (tDCS)

The tDCS parameters were as previously described (chapter 5; O'Shea *et al.* 2017). Stimulation was delivered by a battery driven DC stimulator (Neuroconn GmbH, Ilmenau, Germany) connected to two 7×5 cm sponge electrodes soaked in a 0.9%

saline solution. The electrodes were positioned immediately before stimulation onset and removed as soon as the stimulation ended. The anode electrode was centred over C3 (5 cm lateral to Cz) corresponding to the left primary motor cortex according to the international 10-20 System (Herwig *et al.* 2003). The reference electrode (cathode) was placed over the right supraorbital ridge. During anodal stimulation, the current intensity was set to 1 mA for 20 minutes with a ramp-up and ramp-down period of 10 seconds prior and after the stimulation respectively. During sham stimulation, the current also ramped up and down for 10 sec but no stimulation was delivered during the 20 minutes. Instead, small current pulses (110 μ A over 15 ms) occurred every 550 ms to simulate the tingling sensation associated with real anodal stimulation. Both experimenters and participants were blinded to the stimulation condition.

6.2.4 MRS data analysis

Metabolites were quantified using LCModel (Provencher 1993; Provencher 2001; Provencher 2012), performed on all spectra within the chemical shift range 0.5 to 4.2 ppm. The model spectra of alanine (Ala), ascorbate/vitamin C (Asc), aspartate (Asp), glycerophosphocholine (GPC), phosphocholine (PCho), creatine (Cr), phosphocreatine (PCr), GABA, glucose (Glc), glutamine (Gln), glutamate (Glu), glutathione (GSH), myo-inositol (myo-Ins), Lactate, N-acetylaspartate (NAA), N-acetylaspartylglutamate (NAAG), phosphoethanolamine (PE), scyllo-inositol (scyllo-Ins) and taurine (Tau) were generated based on previously reported chemical shifts and coupling constants by VeSPA Project (Versatile Stimulation, Pulses and Analysis). The unsuppressed water signal acquired from the volume of interest was used to remove eddy current effects and to reconstruct the phased array spectra (Natt *et al.* 2005). Single scan spectra were corrected for frequency and phase variations induced by subject motion before summation. The main difference between 7T and 3T MRS is that glutamine and glutamate cannot be distinguished

with the latter. Instead we used glutamix (Glx), a composite measure of Glu and Gln.

	M1 voxel	V1 voxel
% Grey matter	25.05 (3.65)	49.66 (4.73)
% White matter	59.17 (5.63)	31.61 (4.57)
% CSF	15.78 (5.24)	18.71 (4.42)
SNR	36.23 (3.28)	39.50 (7.65)
FWHM	0.04 (0.007)	0.04 (0.007)
GABA %CRLB	23.59 (4.07)	27.55 (22.27)
Glx %CRLB	5.32 (0.89)	5.10 (2.31)
Cr %CRLB	9.45 (1.68)	10.60 (4.62)
PCr %CRLB	16.64 (3.47)	14.45 (5.40)

Table 6.2: MRS characteristics. Group mean (standard deviation in parenthesis) voxel composition, spectrum signal-to-noise ratio (SNR), full width at half maximum (FWHM), and GABA, Glutamix (Glx), Creatine (Cr) and Phosphocreatine (PCr) relative Cramér-Rao lower bound (%CRLB) are shown for each MRS voxel (M1 and V1).

Table 6.2 shows some quality information on the MRS measurements. Similar to chapter 5, we used the quality filtering method described in chapter 3 (see section 3.4.4.2) to avoid biasing the sample towards high concentration estimates (Kreis 2016). In brief, it involved, for each individual dataset, the computation of an expected relative Cramér-Rao Lower Bound (CRLB) given the concentration estimate and assuming a constant level of noise across all measurements. Datasets for which the Pearson residual between the expected and observed relative CRLB exceeded 2 were excluded from subsequent analysis. Using this quality filtering criterion for GABA, Glx and total Creatine (creatinine + phosphocreatine, labelled tCr), four V1 MRS datasets (all age-matched healthy individuals) were discarded and no M1 MRS dataset was discarded.

Tissue correction is an important step in the analysis of MRS data. This is particularly true for cohorts of elderly individuals because of the atrophy that accompanies normal ageing (Porges *et al.* 2017b). The output of LCmodel represents the metabo-




Figure 6.1: Magnetic Resonance Spectroscopy data quality. **A.** Spectroscopy voxels group overlap map. The M1 voxel was centred on the left central sulcus in 22 participants (in green, $z = 52$). The control V1 voxel was placed centred on the bilateral calcarine sulcus in 20 participants (in blue, $z = 2$). Colour code represents the degree of overlap. All images are displayed in radiological convention (i.e. left side of the image corresponds to the right side of the brain). **B.** Example raw MRS spectrum and LC model fit from one participant. The fitted LCModel (in red) is plotted overlaid on the raw data (in black). The difference between the data and model (residuals) is shown at the top and the baseline is shown at the bottom. **C.** Inter-individual partial correlation between age and total Creatine (tCr) absolute concentration estimate, controlling for the fraction of WM and GM in the M1 voxel (in green) and the V1 voxel (in blue). In the M1 voxel, absolute tCr was negatively correlated with age ($r_{(22)} = -0.46, p = 0.04$). This was not the case in the V1 voxel ($r_{(18)} = -0.06, p = 0.81$).

lite concentration in the entire volume of interest. Therefore, if the fraction of neural tissue containing the metabolite of interest decreases, due to atrophy (Good *et al.* 2002), the concentration of this metabolite in the MRS voxel will necessarily be lower. However, this depletion does not reflect a reduction in the metabolite concentration within the tissue, but rather is a by-product of atrophy. To account for this possible confound, it is crucial that the analysis controls for the tissue content in the MRS voxel. There is currently no consensus in the MRS community on the best method to do so. Most correction techniques rely on implicit assumptions about the distribution of the metabolite of interest within the different tissue compartments (Harris *et al.* 2015; Porges *et al.* 2017a). However, this distribution might change with ageing, making this type of correction potentially invalid. To avoid this caveat, all analyses reported in this chapter used non-tissue corrected concentration estimates and in-

cluded the percentage of grey matter (GM) and white matter (WM) present in the voxel as confound variables (method used in Scholl *et al.* 2017). This partial volume correction approach makes no assumption regarding the distribution of GABA and Glx within the different tissue types, which makes it more valid in the context of ageing. FMRIB’s automated segmentation tool (Patenaude *et al.* 2011), part of the FSL software library (Smith *et al.* 2004), was used on the T1-weighted images to calculate the percentage of grey matter (GM), white matter (WM) and cerebrospinal fluid (CSF) present in the volume of interest (labelled [GM], [WM] and [CSF] respectively).

Across individuals, the tCr concentration estimate was negatively correlated with age in the M1 voxel ($r_{(22)} = -0.46, p = 0.04$; [Figure 6.1C](#)) but not in the V1 voxel ($r_{(18)} = -0.06, p = 0.81$; [Figure 6.1C](#)). The relationship with age in the M1 voxel may be problematic when using tCr for internal referencing (see [chapter 3](#)). Therefore, we report absolute concentration estimates rather than ratios to tCr for GABA and Glx in the main text of this chapter. [Appendix C](#) reports the same analyses with ratios to tCr rather than absolute concentrations. This did not affect the findings reported in this chapter.

6.2.5 Computational analysis of prism adaptation data

The weighted two-state linear time-invariant model described in [chapter 4](#) was fitted to the adaptation data. On CLP trials, the model is defined as follows:

$$\begin{aligned}
 e(n) &= bias_{CLP} + d - X(n) \\
 X(n) &= x_f(n) + x_s(n) \\
 x_f(n) &= A_f x_f(n-1) + B_f e(n-1) \\
 x_s(n) &= A_s x_s(n-1) + B_s e(n-1) \\
 0 < A_f < A_s < 1; \quad 0 < B_s < B_f < 1
 \end{aligned}
 \tag{6.1}$$

where the endpoint error on the n -th trial $e(n)$ is modelled as the sum of a bias term $bias_{CLP}$ and a prism effect d from which is removed the sum of the states of the two systems $X(n)$. The learning and retention rates (B and A) determine the speed of learning and decay of each system. On OLP trials, where prisms are removed and no visual feedback is available, the weighted two-state model becomes:

$$\begin{aligned}
 e(n) &= bias_{OLP} - X(n) \\
 X(n) &= w_f x_f(n) + w_s x_s(n) \\
 x_f(n) &= A_f x_f(n-1) \\
 x_s(n) &= A_s x_s(n-1) \\
 0 < A_f < A_s < 1; \quad 0 < B_s < B_f < 1 \\
 0 < w_f < 5; \quad 0 < w_s < 5
 \end{aligned} \tag{6.2}$$

Equations 6.1 and 8.6 differ in four ways. First, a different bias term is used because of baseline differences in pointing accuracy between CLP and OLP. Second, the prism effect is set to zero to reflect removal of the prisms. Third, the learning rates of both systems (B_f , B_s) are set to zero because the absence of visual feedback prevents error-dependent learning on OLP. Fourth, weighting factors (w_f , w_s) are introduced to allow for a differential contribution of the fast and slow systems to motor behaviour during AE measurements (see chapter 4).

The fitting procedure was as described in chapters 4 and 5, i.e. we used a least-mean-squared-error procedure implemented in MATLAB (The MathWorks Inc., version R2014b). The six free parameters (A_f , A_s , B_f , B_s , w_f , w_s) were first fitted to the group mean data. When fitting individual datasets, the learning and retention rates of the fast and slow systems were set to the group values and only w_f and w_s were estimated on an individual basis. This approach reduced the number of free pa-

rameters, enabling individuals' AEs to be modelled solely in terms of the two weights, one per system. Within this modelling framework, temporal credit assignment (i.e. the computation of the timescale over which to retain the solution of the adaptation, Körding *et al.* 2007) can be conceptualised as the relative weight assigned to the slow versus fast system, computed as: $R_w = w_s/(w_f + w_s)$. As shown in chapter 5 (Figure 5.6), higher R_w predicts more stable AE in young adults. Here, we investigated whether the same relationship holds in a sample of older participants.

6.2.6 Statistical analysis

Unless specified otherwise, all statistical tests reported are two-tailed, and have been performed on the assumption of normality (Shapiro-Wilk confirmed). Non-parametric testing was used if this assumption was violated.

Correlation analyses were performed in MATLAB (The MathWorks inc. version R2014b) using Pearson correlation coefficients. Because GABA is synthesised from glutamate, the concentrations of the two neurotransmitters are typically correlated in the brain (M1 GABA $\times$ M1 Glx: $r_{(22)} = 0.34$, $p = 0.13$; V1 GABA $\times$ V1 Glx: $r_{(16)} = 0.16$, $p = 0.55$). It is therefore crucial to orthogonalise all effects with respect to the other neurotransmitter (Stagg *et al.* 2011a; Jocham *et al.* 2012), i.e. it is necessary to compute partial correlations. Thus, when analysing the relationship between the absolute concentration in GABA within a voxel and any other variable, the concentration of the other neurotransmitter (Glx) was regressed out (as well as the fraction in GM and WM, see section 6.2.4). Correlation coefficients were compared using Fisher's r-to-z transformation (Fisher 1915).

In order to fully characterise the mechanistic links underlying the observed correlations between age, M1 GABA, and AE 24 hours after PA (see Results), a mediation

analysis was used. Mediation was conducted using regression with bootstrapping to ascertain whether the level of M1 tonic inhibition accounts for the link between age and long-term retention of the AE. This analysis was run using the PROCESS macro (<http://www.processmacro.org/index.html>; Hayes 2017) for the SPSS Statistics software package (IBM, Armonk, NY, USA). It included age as the independent variable (X), M1 GABA and Glx absolute concentrations as mediators (M_1 , M_2), and the block mean error on OLP 24 hours after PA normalised by the baseline (pre-PA) deviation as the dependent variable (Y), and controlled for the fraction of GM and WM in the M1 voxel (C_1 , C_2). The Sobel test was used to determine the significance of indirect pathways (Sobel 1982). The percentage mediation (P_M) was calculated as the fraction of total effect (c) accounted by indirect effects (ab_1 or ab_2).

6.3 Results

6.3.1 Comparison of prism adaptation data between young and older adults

Elderly participants ($n = 27$) underwent 100 trials of prism exposure spaced across blocks, interleaved with blocks of AE measurement (Figure 6.2A). Their CLP and OLP performance during the adaptation phase were first compared to those of young healthy adults tested in the same conditions in the experiment reported in chapter 5. Compared to older participants, young healthy adults performed faster movements on CLP, a type of pointing in which speed was emphasised (young adults: mean CLP movement speed = 318.48 ms, SEM = 11.06; older adults: mean CLP movement speed = 422.30 ms, SEM = 19.06; contrast: $t_{(49)} = -4.56$, $p < 10^{-4}$). However, no such difference was detected on OLP, a type of pointing that did not require speeded movements (young adults: mean OLP movement speed = 806.23 ms, SEM = 37.50; older adults: mean OLP movement speed = 828.96 ms, SEM = 25.12; contrast:

$t_{(49)} = -0.51, p = 0.61$).

Two separate mixed analyses of variance (ANOVAs) for CLP blocks (E1-6) and OLP blocks (AE1-6) were performed on the mean block error data with the within-individual factor of Block (1-6) and between-individual factor of Group (young versus older). Participants from both age groups demonstrated the ability to gradually correct pointing accuracy during prism exposure (Figure 6.2A: E1-6; Main effect of Block: $F(5, 245) = 130.10, p < 10^{-7}$). As a result, both age groups significantly shifted leftward on OLP, demonstrating the gradual development of sensorimotor AEs (Figure 6.2A: AE1-6; Main effect of Block: $F(5, 245) = 26.73, p < 10^{-7}$).




Figure 6.2: Prism adaptation and retention data. A. The x-axis represents trial number and the y-axis pointing accuracy relative to baseline (angular error, in degrees, rightward errors are positive). Data (group mean, error bar = SEM) from young participants (in black, $n = 24$, mean age = 25) and older participants (in red, $n = 27$, mean age = 68) are shown. Black wedges indicate blocks throughout which 10° right-shifting prisms are worn. Participants received visual feedback on CLP trials (shaded dark grey), allowing error-dependent learning to occur. By contrast, no visual error feedback was provided on OLP (shaded light grey), allowing the measurement of prism AE. Both age groups demonstrated significant error reduction on CLP and leftward shift on OLP (AE). There was a trend towards greater AEs in older participants during PA ($F(1, 49) = 3.35, p = 0.07$). At both retention time points (10 minutes and 24 hours), older adults demonstrated significant retention of the AE (asterisk: $p < 0.05$).

We found no evidence of an impairment in error reduction in older adults. In fact, there was a small trend towards generally more accurate CLP in older participants (Figure 6.2A: E1-6; Main effect of Group: $F(1, 49) = 1.83$, $p = 0.18$). However, this small trend disappeared when controlling for individuals' mean movement speed on CLP ($F(1, 48) = 0.64$, $p = 0.43$), indicating that this apparently better accuracy in older adults was likely to be an artefact resulting from greater in-flight movement correction due to slower movements. Older participants had slightly greater AEs throughout the adaptation phase (Figure 6.2A: AE1-6; Main effect of Group: $F(1, 49) = 3.35$, $p = 0.07$). This effect was comparable when controlling for individuals' mean movement speed on OLP ($F(1, 48) = 3.19$, $p = 0.08$), indicating that speed of movement during AE measurements had no influence on this effect. No interaction was significant.




Figure 6.3: Model fit in young and older participants. A-B. The x-axis represents trial number and the y-axis pointing accuracy relative to baseline (angular error, in degrees, rightward errors are positive). Graphs show normalised PA data (in black, group mean, error bar = SEM) overlaid with the weighted two-state model fit (in red) and the trial-by-trial state of the fast (in green) and slow (in blue) systems for young (A) and older (B) participants. **C.** Graph shows the relative weight in both age groups. Individual estimates (black dots) are overlaid over the group mean (red line), a 95% confidence interval box (in red) and 1SD box (in blue). There was no significant difference between the two groups ($Z = 1.37$, $p = 0.17$).

To compare the hidden internal states driving PA between young and older adults, the weighted two-state model was fitted to the group mean data of both age groups (Figure 6.3). At the group level, the learning and forgetting rates of the two systems were comparable across the two age groups, apart from a slightly increased learning rate of the fast system in older adults (young adults: $B_f = 0.13$; older adults: $B_f = 0.17$; Figure 6.3), reflecting faster error correction, as reported above. However,

the two groups seemed to differ in terms of the weight assigned to the two systems on AE measures. Compared to young healthy adults, the fit of the group mean data of older participants yielded a lower weight of the fast system (young adults: $w_f = 1.11$; older adults: $w_f = 1.01$; [Figure 6.3](#)) and a greater weight of the slow system (young adults: $w_s = 0.82$; older adults: $w_s = 0.94$; [Figure 6.3](#)), resulting in a greater relative weight (young adults: $R_w = 0.42$; older adults: $R_w = 0.48$; [Figure 6.3](#)). To confirm whether the PA profile of older participants was indeed characterised by a greater relative weight (i.e. greater relative contribution of the slow system to the overall AEs), the weighted two-state model was fitted to individual PA data of the two groups. When doing so, the learning and retention rates were set to the values estimated at the group level for each age group respectively, and only the weights of the two systems were estimated. Non-parametric testing (Mann–Whitney–Wilcoxon test) was used to compare the relative weight between the two age groups because the assumption of normality was violated. No significant difference was found between the two age groups ($Z = 1.37$, $p = 0.17$; [Figure 6.3C](#)), indicating that the model could not capture computational differences between young and older participants.

In young healthy adults, a greater relative weight predicted a greater and more stable AE memory trace at retention ([chapter 5: Figure 5.6](#)). Here, we reproduced the identical analysis to test whether that was also the case in older adults. A robust multiple linear regression was used to test whether the initial magnitude (average of the first three OLP trials of the block) and linear slope of the AE at retention was associated with R_w . In older participants, this analysis failed to identify a link between R_w and the property of the AE 10 minutes (magnitude: $\beta = -0.10$, $p = 0.67$; stability: $\beta = -0.30$, $p = 0.22$) or 24 hours after PA (magnitude: $\beta = 0.32$, $p = 0.18$; stability: $\beta = 0.02$, $p = 0.95$).

In summary, the weighted two-state model provided a good fit to motor behaviour during PA in older adults ([Figure 6.3](#)) but did not relate quantitatively to the AE memory trace at retention. This suggests that in this population, there may be additional processes (potentially occurring after PA) influencing the AE memory trace than those captured by the weighted two-state model. Nonetheless, we asked in the following section whether, as in young adults ([chapter 5](#)), sensorimotor cortex excitation:inhibition was associated with the relative weight in older participants.

6.3.2 Relationship between excitation:inhibition and relative weight in older adults

In young adults ([chapter 5](#)), greater basal M1 excitation:inhibition ratio was associated with greater relative weight ([Figure 5.6](#)). Here, we reproduced the same analysis to test whether the same applied in a sample of older adults. The only difference was that, due to lower magnetic field strength of the MRI scanner, the E:I ratio was quantified as the relative concentration of glutamix-to-GABA (Glx:GABA) rather than glutamate-to-GABA as in the previous chapter.

Robust linear regression did not confirm the expected positive relationship in older adults. Instead, the opposite relationship was found. In older adults, greater M1 E:I ratio was associated with greater relative contribution of the fast system to the AE during the adaptation (M1 Glx:GABA $\times R_w$: $\beta = -0.51$, $p = 0.03$; [Figure 6.4](#)). This relationship was still significant after controlling for age ($p < 0.05$). However, the relevance of this result for behavioural retention of the AE is unclear given the absence of predictive value of the relative weight beyond the adaptation phase (i.e. no association with the AE 10 minutes or 24 hours after PA; [section 6.3.1](#)).




Figure 6.4: Relationship between M1 excitation:inhibition and relative weight. Across individuals, greater M1 excitation:inhibition was associated with greater relative contribution of the fast system to AEs during PA ($\beta = -0.51$, $p = 0.03$). The directionality of this relationship is inverted compared to young adults ([chapter 5](#)). Colour code represents the weight of individual data points in the robust linear regression.

Given the outlined limited ability of the weighted two-state model to extract computational parameters that are relevant for long-term retention in older adults, the next sessions report analyses of the raw behavioural data at retention.

6.3.3 Older adults show more persistent after-effects

Retention data could not be compared across the two age groups because young participants performed 20 extra CLP trials during the adaptation phase to account for their faster pace ([chapter 5](#)). Subsequent analyses therefore tested for cross-sectional correlations between age and AE 10 minutes and 24 hours after PA in older adults only. Significant AEs were detected at both time points (10 minutes: mean $AE_{10min} = -4.52^\circ$, $SEM = 0.41$; $t_{(26)} = -11.10$, $p < 10^{-10}$; 24 hours: mean $AE_{24hrs} = -1.17^\circ$, $SEM = 0.53$; $t_{(26)} = -2.22$, $p = 0.04$; [Figure 6.2A](#)). Age showed no correlation with the magnitude of AE 10 minutes after PA (Age $\times$ AE_{10min} : $r_{(27)} = -0.03$, $p = 0.89$; [Figure 6.5A](#)). By contrast, age was significantly correlated with the AE 24 hours after the adaptation (Age $\times$ AE_{24hrs} : $r_{(27)} = -0.48$, $p = 0.01$; [Figure 6.5B](#)). In this population, older age was associated with larger AE (i.e. more left-shifted OLP accuracy) 24 hours after PA. This correlation was unchanged when controlling for the

magnitude of AE obtained at the end of the adaptation or 10 minutes after PA (both $p < 0.05$), indicating that it reflects increased persistence of the AE in older adults rather than generally greater AEs throughout the experiment.




Figure 6.5: Effect of age on AE persistence. **A.** Graph shows the correlation between age (x-axis) and block mean prism AE 10 minutes after PA (y-axis). This correlation was not significant ($r_{(27)} = -0.03$, $p = 0.89$). **B.** Graph shows the correlation between age (x-axis) and block mean AE 24 hours after PA (y-axis). Across individuals, older age was associated with greater (i.e. more negative) AE 24 hours after PA, indicating a more persistent memory trace.


A *post hoc* median split of the data revealed that only the oldest half of the participants (age > 67.8 , $n = 13$) showed significant AEs the next day (mean $AE_{24hrs} = -2.35^\circ$, $SEM = 0.78$, $p = 0.01$), whereas the youngest half did not (mean $AE_{24hrs} = 0.17^\circ$, $SEM = 0.60$, $p = 0.77$).

6.3.4 Older adults have lower GABAergic inhibition

Several studies have associated ageing with a general down-regulation of GABAergic inhibition (Peinemann *et al.* 2001; Oliviero *et al.* 2006; Gao *et al.* 2013; Heise *et al.* 2013; He *et al.* 2016; Porges *et al.* 2017a; Porges *et al.* 2017b). In our sample, partial correlation analyses (controlling for [GM], [WM] and Glx concentration) confirmed this observation in M1 (Age $\times$ M1 GABA: $r_{(22)} = -0.52$, $p = 0.01$, 1-

tail; Figure 6.6A) and there was a similar trend in V1 (Age $\times$ V1 GABA: $r_{(16)} = -0.43$, $p = 0.07$, 1-tail; Figure 6.6A). Neither region of interest showed any correlation between age and Glx levels (Age $\times$ M1 Glx: $r_{(22)} = -0.16$, $p = 0.51$, Age $\times$ V1 Glx: $r_{(16)} = 0.07$, $p = 0.83$; Figure 6.6). The expression of GABA and Glx levels as ratios to tCr did not change these results (Appendix C).

A) M1 voxel



B) V1 voxel




Figure 6.6: Effect of age on spectroscopy measurements. Graphs show the partial correlation between age (x-axis) and GABA or Glx absolute concentration in M1 (A.) and V1 (B.) MRS voxels. The fraction of grey matter and white matter present in the voxel were regressed out from these relationships. Additionally, the effect of age on neurotransmitter concentration in each voxel is orthogonalised with respect to the other neurotransmitter. Across participants, older age was associated with lower GABA concentration in M1 ($r_{(22)} = -0.52$, $p = 0.01$, 1-tail) and V1 ($r_{(16)} = -0.43$, $p = 0.07$, 1-tail). This relationship was not found for Glx (both $p > 0.50$).

Note that this cross-sectional age-related decrease in GABAergic inhibitory tone in the context of relatively preserved glutamatergic excitation resulted in increased M1 excitation:inhibition ratio (computed as Glx:GABA) with older age (Age $\times$ M1

Glx:GABA: $r_{(22)} = 0.44$, $p = 0.03$, 1-tail). A similar trend was found in the V1 voxel (Age $\times$ V1 Glx:GABA: $r_{(16)} = 0.39$, $p = 0.09$, 1-tail).

6.3.5 Lower M1 GABA is associated with more persistent after-effect




Figure 6.7: Effect of M1 neurotransmitters on AE persistence. Graphs show the partial correlation between GABA or Glx absolute concentration in M1 (A.) and V1 (B.) and prism after-effect (AE) 24 hours after PA, a marker of AE persistence. The fraction of grey matter and white matter present in the voxel were regressed out from these relationships. Additionally, the effect of age on neurotransmitter concentration in each voxel is orthogonalised with respect to the other neurotransmitter. As predicted, individuals with lower M1 GABA concentration showed more persistent AEs (i.e. more left-shifted OLP behaviour on day 2) (A, $r_{(22)} = 0.74$, $p < 10^{-3}$, 1-tail). This relationship was specific to M1 and not V1 (V1 GABA: $r_{(16)} = -0.06$, $p = 0.85$; contrast: $Z = 2.79$, $p = 0.005$), and to GABA and not Glx (M1 Glx: $r_{(22)} = -0.03$, $p = 0.90$, contrast: $Z = 2.84$, $p = 0.005$).

Next, we tested our key hypothesis that inter-individual differences in M1 but not V1 GABA levels should co-vary with inter-individual differences in AE persistence. More precisely, based on findings from the previous chapter ([chapter 5](#)), we predicted that, across individuals, lower M1 GABA levels should be associated with more persistent AEs, i.e. more left-shifted OLP on day 2. A partial correlation analysis (controlling for $[GM]_{M1}$, $[WM]_{M1}$ and M1 Glx level) confirmed this prediction (M1 GABA $\times$ AE_{24hrs}: $r_{(22)} = 0.74$, $p < 10^{-3}$, 1-tail; [Figure 6.7A](#)). The anatomical specificity of this relationship was demonstrated by an absence of such correlation in the V1 voxel (V1 GABA $\times$ AE_{24hrs}: $r_{(16)} = -0.06$, $p = 0.85$; contrast: $Z = 2.79$, $p = 0.005$; [Figure 6.7A](#)). Its neurochemical specificity was demonstrated by an absence of correlation between M1 Glx and AE 24 hours after PA (M1 Glx $\times$ AE_{24hrs}: $r_{(22)} = -0.03$, $p = 0.90$; contrast: $Z = 2.84$, $p = 0.005$; [Figure 6.7A](#)). V1 Glx also did not show any association with long-term retention of the AE (V1 Glx $\times$ AE_{24hrs}: $r_{(16)} = 0.22$, $p = 0.48$; [Figure 6.7B](#)). The same pattern of results was found when expressing metabolite concentrations as ratios to tCr ([Appendix C](#)).


6.3.6 Mediation analysis

In order to more fully investigate the potential role of M1 GABAergic inhibition as a mechanistic link between age and AE persistence, a mediation analysis was conducted. The goal was to test whether the reduced M1 GABA levels observed in older adults ([Figure 6.6A](#)) explained the increase in AE persistence ([Figure 6.5B](#)). Mediation approaches assess the mechanism by which two variables (in this case, age and AE_{24hrs}) are related. It is well suited to a situation in which candidate mediators M_i (M1 Glx, M1 GABA) are the logical effect of an independent variable X (age) and cause of a dependent variable Y (AE_{24hrs}).

The full results of the mediation analysis are shown in [Table 6.3](#). There was

strong evidence that M1 GABA mediated the link between age and long-term retention (AE_{24hrs}), with a percentage mediation of $P_M = 65.98\%$. Age was no longer a significant predictor of AE 24 hours after PA when controlling for M1 neurochemistry ($c' = -0.28$, $p = 0.34$), consistent with full mediation. The indirect effects were tested and contrasted against one another using a bootstrap estimation approach with 5000 samples (Shrout & Bolger 2002). These results indicated a significant mediation by M1 GABA ($ab_1 = -0.50$, $SE = 0.24$, $Z = -2.06$, $p = 0.04$) but not by M1 Glx ($ab_2 = 0.02$, $SE = 0.09$, $Z = 0.21$, $p = 0.83$; contrast: $p = 0.04$).

In summary, the mediation analysis indicates that the greater long-term retention of the AE observed in older adults is a consequence of their lower M1 GABAergic inhibition. Fully investigating the causal nature of the relationships between age, M1 GABA and AE persistence would require a longitudinal study, and is outside the focus of this thesis. What matters the most in the context of the present study, is the preserved association between M1 tonic GABAergic inhibition and AE persistence in older adults.



Antecedent		Consequent															
		M ₁ (GABA)			M ₂ (Glx)			Y (AE _{24hrs})									
		<i>Coeff.</i>	<i>SE</i>	<i>p</i>		<i>Coeff.</i>	<i>SE</i>	<i>p</i>		<i>Coeff.</i>	<i>SE</i>	<i>p</i>		<i>Coeff.</i>	<i>SE</i>	<i>p</i>	
X (age)	a₁	-0.78	0.28	<i>0.01</i>	a₂	-0.31	0.27	<i>0.27</i>	c	-0.76	0.29	<i>0.02</i>	c'	-0.28	0.28	<i>0.34</i>	
M ₁ (GABA)		-	-	-		-	-	-		-	-	-	b₁	0.65	0.20	<i>0.004</i>	
M ₂ (Glx)		-	-	-		-	-	-		-	-	-	b₂	-0.06	0.20	<i>0.78</i>	
C ₁ (GM)		-0.68	0.31	<i>0.04</i>		-0.31	0.30	<i>0.33</i>		-0.4	0.32	<i>0.22</i>		0.02	0.29	<i>0.94</i>	
C ₂ (WM)		-0.53	0.23	<i>0.03</i>		-0.70	0.22	<i>0.006</i>		-0.32	0.23	<i>0.19</i>		-0.01	0.25	<i>0.96</i>	
		R ² = 0.34				R ² =0.38				R ² = 0.33				R ² = 0.60			
		F(3,18) = 3.16				F(3,18) = 3.72				F(3,18) = 2.99				F(5,16) = 4.88			
		<i>p</i> = 0.05				<i>p</i> = 0.03				<i>p</i> = 0.06				<i>p</i> = 0.007			

Table 6.3: M1 GABA mediates the relationship between age and long-term sensorimotor memory. The mediating effect of tonic M1 GABA in the relationship between age and prism AE 24 hours after PA was significant ($ab_1 = -0.50$, $SE = 0.24$, $Z = -2.06$, $p = 0.04$) but M1 Glx had no such mediating effect ($ab_2 = 0.02$, $SE = 0.09$, $Z = 0.21$, $p = 0.83$; contrast: $p = 0.04$). Asterisks and red arrows indicate significant paths ($p < 0.05$). Full details of the statistics are provided in the table.

6.3.7 Effect of online M1 a-tDCS on PA in elderly adults

A follow-up tDCS experiment was run to evaluate the causal status of the relationship between M1 GABAergic inhibition and adaptation memory reported above (Figure 6.4, 6.7 & 6.3). Data collection for this experiment is still going (target: $n = 24$, to match the experiment reported in chapter 5) so the results reported here are preliminary. Eighteen older participants underwent the same experimental procedure (PA on day 1, retention on day 2) but received left M1 sham or anodal tDCS during PA, in counterbalanced order. This intervention is known to reduce M1 GABA concentration in older adults (Antonenko *et al.* 2017). Based on the previous findings reported in this chapter, we had three predictions:

1. Across older participants, lower excitation:inhibition was associated with lower relative weight (Figure 6.4). If this basal relationship is causal, then a-tDCS should result in a 'paradoxical' decrease in relative weight, characterised by greater AEs early during PA but similar AEs towards the end of the adaptation.
2. Nonetheless, because of the negative correlation between M1 GABA and AE persistence (Figure 6.7), M1 a-tDCS should enhance the magnitude of AE present 24 hours after PA. This might not be the case 10 minutes after PA because of the absence of relationship with M1 GABA at this retention time point.
3. The down-regulation of GABAergic inhibition observed in older participants might weaken the effect of anodal tDCS by imposing a ceiling effect. Individuals with higher baseline M1 excitation:inhibition ratio are therefore likely to show little to no response to M1 anodal tDCS.





A) Behavioural results**B) Relative weight****C) Retention data****D) M1 E:I x tdcS effect**

Figure 6.8: Effect of M1 a-tDCS during prism adaptation. **A.** The x-axis represents trial number and the y-axis pointing accuracy in reference to baseline accuracy (i.e. change from baseline). Pointing accuracy (in degrees) is plotted when anodal (red) or sham (blue) was applied to the left M1 during prism adaptation (group mean $\pm$ 1 SEM). Black wedges indicate blocks throughout which prisms were worn (CLP trials). During Adaptation (prism exposure, E1-7), participants saw the outcome of each trial, allowing error-dependent learning to proceed. The after-effect was measured without visual feedback (AE1-7, shaded light grey). Retention was measured after 10 minutes of blindfolded rest, and after 24 hours. **B.** As predicted, there was a trend towards a stimulation induced reduction in relative weight ($t_{(17)} = -0.94$, $p = 0.18$, 1-tail). **C.** As predicted, M1 a-tDCS during PA did not affect the magnitude of AE 10 minutes after PA ($t_{(17)} = 0.04$, $p = 0.97$), but there was a trend towards increased AEs 24 hours after PA ($t_{(17)} = -1.12$, $p = 0.14$, 1-tail). Graphs (B-C) show the group mean (line), a 95% confidence interval box (dark colour) and 1SD box (light colour). Individual data are plotted as a spaghetti plot overlaid on the box plot. **D.** Individuals with lower excitation:inhibition (Glx:GABA) in M1 showed a greater response to anodal tDCS (i.e. more negative AE 24 hours after PA with a-tDCS compared to sham tDCS; $r_{(14)} = 0.54$, $p = 0.02$, 1-tail). A: left M1 anodal tDCS; S: left M1 sham tDCS. A = anodal tDCS condition; S = sham tDCS condition.

Preliminary results provided support in favour of all three predictions (Figure 6.8).

There was a trend towards a decrease in relative weight during PA combined with anodal tDCS compared to sham tDCS ($\Delta R_w = -0.04$, SEM = 0.05, $t_{(17)} = -0.94$, $p =$

0.18, 1-tail; [Figure 6.8B](#)). The magnitude of AE obtained 10 minutes after PA was unchanged in the a-tDCS condition ($\Delta\text{AE}_{10\text{min}} = 0.02$, $\text{SEM} = 0.61$, $t_{(17)} = 0.04$, $p = 0.97$; [Figure 6.8C](#)). However, compared to the sham condition, there was a trend towards increased (i.e. more negative) AEs 24 hours after PA in the anodal condition ($\Delta\text{AE}_{24\text{hrs}} = -0.98$, $\text{SEM} = 0.87$, $t_{(17)} = -1.12$, $p = 0.14$, 1-tail; [Figure 6.8C](#)). Finally, across individuals, the magnitude of the latter effect was influenced by baseline M1 excitation:inhibition. Participants with higher excitation:inhibition in M1 showed a smaller or even opposite response to a-tDCS ($\text{M1 Glx:GABA} \times \Delta\text{AE}_{24\text{hrs}}$: $r_{(14)} = 0.54$, $p = 0.02$, 1-tail; [Figure 6.8D](#)). This correlation was still significant when controlling for age ($r_{(14)} = 0.52$, $p = 0.03$), and was not found for V1 excitation:inhibition ratio ($\text{V1 Glx:GABA} \times \Delta\text{AE}_{24\text{hrs}}$: $r_{(10)} = -0.008$, $p = 0.98$; contrast: $Z = 1.27$, $p = 0.10$, 1-tail).

6.3.8 Neurotransmitter levels in neglect

A secondary goal of this chapter was to address the same questions as above in the clinical population of interest, i.e. post-stroke neglect patients. However, because of recruitment challenges discussed in [chapter 8](#), only six patients were recruited, five of whom underwent MR spectroscopy. This sample size was too small for correlation analyses so we report here the only analysis that could be done: the comparison of Glx and GABA concentrations between neglect patients and age-matched healthy controls. This analysis was run using non-parametric permutation testing implemented in PALM (PALM version alpha-109; <https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/PALM>) (Winkler *et al.* 2014; Winkler *et al.* 2016a; Winkler *et al.* 2016b), controlling for the fraction of WM and GM present in the MRS voxel.

The results indicate no significant difference in GABA, Glx, or Glx:GABA in either voxel (all $p > 0.30$; [Figure 6.9](#)), except for greater Glx concentration in the visual voxel




Figure 6.9: Neurotransmitter levels in neglect. Graphs show GABA, Glx and Glx:GABA in the M1 voxel (A) and V1 voxel (B) for neglect patients (N, in red) and age-matched healthy controls (C, in blue). There was no difference in any of the measures except for greater Glx concentration in the V1 voxel in neglect patients compared to age-matched healthy controls ($p = 0.001$, significant with a Bonferroni threshold of $p < 0.008$). On each box, the central mark indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme data points not considered outliers. Individual data points are overlaid on the box plots. N: neglect patients; C: age-matched control individuals.

in neglect patients compared to age-matched healthy controls ($p = 0.001$, significant with a Bonferroni threshold of $p < 0.008$; Figure 6.9B). This difference might reflect higher excitation in the left visual cortex due to altered inter-hemispheric balance (Kinsbourne 1976; Heilman & Van Den Abell 1980; Corbetta *et al.* 2005). This point will be discussed in more details in chapter 7.

6.4 Discussion

This study aimed to test whether, like in young healthy adults, M1 tonic excitation:inhibition ratio determines retention of prism after-effects in older adults. To test this, MRS measurements of neurotransmitter concentrations were taken at rest

in the left sensorimotor cortex and bilateral occipital cortex. On a separate session, participants performed PA, and AEs were probed at an early (10-minutes) and late (24-hours) retention time point. Our results indicate that older adults show increased long-term retention of the AE due to lower M1 tonic GABAergic inhibition (Table 6.3). This is compatible with the known role of M1 excitation:inhibition in the formation of adaptation memory (chapter 5; O'Shea *et al.* 2017). In a second experiment, anodal or sham tDCS was applied to M1 during PA to test the main prediction that online excitatory stimulation of the sensorimotor network would result in an increase in long-term retention of the AE. To date, only 18 participants completed the study (i.e. two thirds of the target sample size) and preliminary results indicate a trend towards a stimulation-induced increase in AE 24 hours after PA. This effect was greater in participants with low basal M1 excitation:inhibition ratio, confirming the role of M1 neurochemistry in sensorimotor memory formation.

The dysregulation of tonic GABAergic signalling in ageing is a well documented phenomenon (Peinemann *et al.* 2001; Kaiser *et al.* 2005; Oliviero *et al.* 2006; Sailasuta *et al.* 2008; Chang *et al.* 2009; Gao *et al.* 2013; Porges *et al.* 2017a; Porges *et al.* 2017b). Typically, perturbation of the excitation:inhibition balance via down-regulation of inhibition has been reported to be detrimental for performance (Heise *et al.* 2013; Levin *et al.* 2014; Rozycka & Liguz-Leczna 2017). However, to our knowledge, the consequences of the age-related decrease in M1 inhibition on sensorimotor adaptation had never been elucidated before. Our data demonstrate a 'paradoxical' beneficial effect of ageing on the retention of the prism AE via a reduction of tonic GABAergic inhibition in M1 (Table 6.3). During adaptation, artificially reducing inhibition in M1 via a-tDCS (Antonenko *et al.* 2017) enhanced long-term retention of the AE, which is similar to what was observed in the first experiment when inhibition was reduced by a natural process of ageing. Crucially, these results indicate

that the role of M1 inhibition in long-term retention of the AE is preserved in older adults (O'Shea *et al.* 2017). It has implications in the broader scheme of this thesis because it suggests that retention of the prism AE could also be optimised in neglect patients, who show similar M1 excitation:inhibition ratio as controls, by locally reducing GABAergic inhibition within M1 during PA.

Young and older adults seem to differ however in terms of the computational correlates of M1 excitation:inhibition balance. In young adults, greater E:I ratio was associated with greater relative contribution of the slow system to the AE during adaptation. The opposite was found in older adults. Greater E:I ratio was instead associated with lower relative contribution of the slow system to the AE during adaptation. However, in this population, the weighted two-state model did not predict retention data, so lower relative weight was not associated with a more labile AE memory trace (as is the case in young participants). This absence of relationship potentially suggests different and/or additional processes taking place during or after PA in older adults, and altering the relationship between PA profile and behavioural retention. Further work is therefore needed to clarify the computational differences during and after PA in older adults (King *et al.* 2013). Some authors suggested that these may involve reduced explicit knowledge about the perturbation (Bock 2005; Hegele & Heuer 2010), increased reliance on erroneous sensorimotor predictions (Wolpe *et al.* 2016) or decline in model-free reinforcement learning (Heuer & Hegele 2014). We suggest that testing modality-specific AEs (e.g. visual straight ahead judgement, proprioceptive straight ahead pointing) and using a computational framework that includes better-defined computational modules in terms of their psychological significance (for example, see the framework described in [Appendix A](#); Haith *et al.* 2009; Wei & Körding 2010; Broucke 2017) might provide some answers to this question.

6.5 Conclusions

The main finding of this study is the preserved relationship between M1 tonic GABAergic inhibition and sensorimotor memory in the ageing brain, hence suggesting a good translatability of our previous findings ([chapter 5](#); O'Shea *et al.* 2017) to neglect patients. This relationship explains why older adults, in whom M1 tonic GABAergic inhibition is reduced, also show enhanced long-term retention of prism after-effects. By considering the neurochemical level, this study shed new light on the age-related physiological alterations underlying changes in sensorimotor adaptation performance. Future work should characterise further the computational basis of such changes.

Chapter 7

Functional and structural basis of spatial neglect

The present chapter investigates the neural bases of chronic spatial neglect in a sample of six patients who participated in the rehabilitation study described in the next chapter ([chapter 8](#)). Consistent with a role of inter-hemispheric balance in the allocation of visual attention in space ([Appendix B](#)), neglect was associated with a rightward shift of attention and a leftward shift of inter-hemispheric dominance in the lateral occipital cortex.

7.1 Introduction

Spatial neglect is a visuospatial attention deficit occurring in approximately one out of two patients at the acute stage following a right hemispheric stroke (Stone *et al.* 1993; Karnath *et al.* 2004; Milner & McIntosh 2005; Verdon *et al.* 2009). It is characterised by the failure to report, respond or orient to novel or meaningful stimuli presented to the contralesional side, when this failure cannot be attributed to either sensory or motor deficits (Heilman & Van Den Abell 1980; Heilman *et al.* 2002). In about half of these patients, the ability to attend to the left side of space spontaneously

recovers during the first few weeks or months following the brain injury (Ringman *et al.* 2004; Milner & McIntosh 2005). However, this leaves a large proportion of right hemispheric stroke patients (approximately 25%) with sustained visuospatial attention deficits that contribute to reducing functional recovery, independence and quality of life (Jehkonen *et al.* 2006; Patel *et al.* 2007; Di Monaco *et al.* 2011). The rehabilitation of these chronic neglect patients is challenging (Barrett *et al.* 2012; Bowen *et al.* 2013), and is the primary clinical goal of this thesis.

The anatomical distribution of lesions that can lead to spatial neglect is remarkably heterogeneous. Decades of lesion mapping studies have linked this syndrome to many cortical and subcortical areas including: 1) parietal regions (superior parietal lobule, angular gyrus, supramarginal gyrus, temporo-parietal junction) (Vallar & Perani 1986; Binder *et al.* 1992; Samuelsson *et al.* 1997; Leibovitch *et al.* 1998; Mort *et al.* 2003; Committeri *et al.* 2006; Verdon *et al.* 2009; Chechlacz *et al.* 2010; Karnath *et al.* 2010; Thiebaut de Schotten *et al.* 2012; Carter *et al.* 2017), 2) temporal regions (superior and middle temporal gyri) (Samuelsson *et al.* 1997; Karnath *et al.* 2001; Karnath *et al.* 2004; Verdon *et al.* 2009; Karnath *et al.* 2010), 3) posterior frontal regions (inferior and middle frontal gyri, precentral gyrus) (Damasio *et al.* 1980; Binder *et al.* 1992; Husain & Kennard 1996; Leibovitch *et al.* 1998; Committeri *et al.* 2006; Verdon *et al.* 2009; Chechlacz *et al.* 2010; Thiebaut de Schotten *et al.* 2012; Carter *et al.* 2017), 4) insular regions (Binder *et al.* 1992; Karnath *et al.* 2004; Chechlacz *et al.* 2010; Carter *et al.* 2017), and 5) basal ganglia (Damasio *et al.* 1980; Vallar & Perani 1986; Karnath *et al.* 2004; Karnath *et al.* 2010). More recently, the interruption of some long-range fronto-parietal and fronto-occipital white matter tracts interconnecting brain areas involved in spatial attention has been proposed to provide a unifying framework for understanding the variety of grey matter loci correlated with spatial neglect (Thiebaut de Schotten *et al.* 2005; Roux *et al.* 2011;

Urbanski *et al.* 2011; Bozzali *et al.* 2012; Thiebaut de Schotten *et al.* 2012; Vallar *et al.* 2014; Lunven *et al.* 2015; Herbet *et al.* 2017). These white matter tracts include the most ventral branches of the superior longitudinal fasciculus (SLF II and SLF III) and the inferior fronto-occipital fasciculus (IFOF). SLF II connects regions of the caudal-inferior parietal cortex (angular gyrus, intraparietal sulcus) to regions of the superior and middle frontal gyrus, and SLF III connects the supramarginal gyrus to the inferior frontal gyrus (Petrides & Pandya 1984; Thiebaut De Schotten *et al.* 2011; Thiebaut de Schotten *et al.* 2011). IFOF connects the ventro-medial occipital cortex to the orbital and polar frontal cortex (Forkel *et al.* 2014). The interruption of these tracts may be the common dysfunction underlying various grey matter lesions associated with spatial neglect.

This hodotopic framework conceptualises neglect as a disconnection syndrome (Doricchi *et al.* 2008; Catani & Thiebaut de Schotten 2012; Thiebaut de Schotten *et al.* 2015) in which the dysfunction of the neural networks involved in the control of visuospatial attention leads to strong spatial biases (He *et al.* 2007; Carter *et al.* 2010; Corbetta & Shulman 2011). It is compatible with a model of visuospatial attention postulated by Corbetta and Shulman that subdivides the cortical centres of spatial attention into a dorsal and a ventral attention network (Corbetta & Shulman 2002; Corbetta *et al.* 2008; Corbetta & Shulman 2011). This model combines elements of the inter-hemispheric competition model (Kinsbourne 1976; Cohen *et al.* 1994) and the hemispatial theory of right hemispheric dominance for visuospatial attention (Heilman & Van Den Abell 1980; Mesulam 1981) ([Appendix B](#)). The dorsal attention network (DAN) includes bilateral intraparietal sulcus/superior parietal lobule (IPS/SPL), frontal eye field (FEF) and lateral occipital cortex (LOC), and mediates top-down voluntary allocation of attention to locations or features (Fox *et al.* 2006; Shulman *et al.* 2009; Shulman *et al.* 2010). The competition between left and right

cortical nodes of the DAN for the allocation of attention to the contralateral side determines the locus of voluntary (top-down) attention (Duecker & Sack 2015; Petit *et al.* 2015). The ventral attention network (VAN) is located predominantly in the right hemisphere and includes the temporoparietal junction (TPJ), anterior insula and ventral frontal cortex. It is involved in the detection of unattended or unexpected stimuli outside the focus of attention, thus triggering automatic (bottom-up) shifts of attention (Fox *et al.* 2006; Shulman *et al.* 2009; Shulman *et al.* 2010). In spatial neglect, the lesion of the VAN, predominantly right lateralised, is thought to provoke a dysfunction in the structurally intact DAN, resulting in a functional imbalance between left (hyperactive) and right (hypoactive) nodes (Corbetta *et al.* 2005; He *et al.* 2007).

Based on this patho-physiological model, many research groups have tried to restore normal parietal inter-hemispheric balance in neglect (Jacquin-Courtois 2015) by intervening physiologically via inhibitory stimulation of the left hyperactive unlesioned hemisphere, and/or excitatory stimulation of the right hypoactive lesioned hemisphere, using TMS (Ko *et al.* 2008; Sparing *et al.* 2009; Sunwoo *et al.* 2013; Yi *et al.* 2016) or tDCS (Oliveri *et al.* 2001; Song *et al.* 2009; Lim *et al.* 2010; Cazzoli *et al.* 2012; Koch *et al.* 2012; Yang *et al.* 2015; Cha & Kim 2016). Unfortunately, the improvements obtained with this approach are restricted to a few visuospatial tests, and the clinical efficacy on functional independence has yet to be proven (Bowen *et al.* 2013).

An alternative, more promising rehabilitation approach is prism adaptation (Chapters 4, 5 & 6), a form of behavioural therapy in which patients gradually learn, through repeated visually-guided movements, to compensate leftward for the presence of a rightward optical shift (Rossetti *et al.* 1998; Serino *et al.* 2009; chapter 1). In ne-

glect patients, this acquired leftward bias (after-effect of prism adaptation) appears to generalise abnormally, leading to improvements in untrained cognitive domains (e.g. visual search, reading, wheelchair driving, mental representation) (for review, see Jacquin-Courtois *et al.* 2013). The induction of a sensorimotor after-effect (leftward shift in open loop pointing or straight ahead pointing) is therefore believed to be an essential prerequisite to cognitive after-effects in neglect (Newport & Schenk 2012). However, very little is understood about the neural correlates of neglect improvement following prism adaptation (Striemer & Danckert 2010). One potential way to address this question is to use brain imaging to test for overlap in the non-lesioned brain regions involved in the patho-physiology of neglect and those underlying prism after-effects. Any shared neural substrate between the two is likely to play a role in mediating neglect recovery after prism adaptation. The first part of this two-step investigation (i.e. the identification of the neural substrates of prism after-effects) was carried in chapters 5 and 6. Hence, the objective of the present chapter is to investigate the neural bases of chronic spatial neglect in a sample of six patients prior to their participation in the rehabilitation study described in the next chapter (chapter 8). Identifying any of the neural circuits mediating neglect and response to prism adaptation would potentially both advance fundamental neuroscience understanding and inform therapeutics.

7.2 Material and methods

7.2.1 Participants

Six patients (1 female, mean age = 61.3 years, SD = 9.4) were recruited from the Oxford Centre for Enablement (OCE, Oxford, UK) and from the community via advertising (for recruitment chart, see Figure 8.1). Recruitment information and characterisation of baseline behavioural performance are detailed in the next chapter

(chapter 8). The present chapter focuses on performance on a test of visuospatial attention (Posner-like task) performed inside the scanner, and on neuroimaging data. Twenty-three age-matched healthy right-handed individuals (0 female, mean age = 67.8, SD = 9.2) without any personal or family history of neurological or psychiatric disease were recruited to provide normative brain imaging data. In total, 27 participants (5 patients and 22 age-matched healthy individuals) had a complete dataset. MR spectroscopy data could not be acquired for the remaining two participants (1 patient and 1 age-matched healthy individual) because of technical failures. All participants gave written informed consent as approved by the Oxford A Research Ethics Committee (REC reference number: 12/SC/0556). No participant reported any side effect from the experimental procedure.

7.2.2 Posner-like task

Inside the scanner, a Posner-like task was used to recruit brain regions involved in visuospatial attention. This task was chosen because 1) it is known to reliably recruit visuospatial attention networks in age-matched healthy individuals and stroke patients (Umarova *et al.* 2009; Umarova *et al.* 2011), and 2) it can be performed relatively rapidly in the scanner, thus maximising tolerability by patients. A mixed block design was used. Task blocks (8 blocks per fMRI run, 38s each, 12 trials per block) were alternated with rest blocks (8 blocks per fMRI run, 16s each). At the beginning of each trial, an arrow was presented centrally for 800 ms. The arrow pointed either to the right or to the left, indicating the location of the upcoming target. Within each block, on two-third of the trials, the target (asterisk) was presented at the cued location for 400 ms after visual offset of the arrow. The remaining one-third of the trials were null events, in which no target appeared after spatial cueing. No invalid trials (i.e. target appears opposite to the cued location) were used so that task-related activity reflected the voluntary orientation of visuospatial attention. Also investigating

brain activity underlying bottom-up stimulus driven shifts of attention would have required an event-related design, which was not feasible within the time constraints of the scanning session. Trials were separated by a fixation cross hair for a jittered period of 1000–2000ms. The task was to report left and right targets as fast as possible by pressing a single button with the right non-paretic and non-neglected index finger. Participants were also instructed to maintain central fixation on the cross hair at all time during the task. Each task block contained 12 randomly intermixed trials, from which 8 were valid (4 events with targets in the right and 4 in the left side) and 4 were null-events (2 with left and 2 with right spatial cueing). All participants underwent a total of 192 trials, distributed over two fMRI runs, each with 8 task blocks. A break of approximately 9 minutes, during which T1-weighted structural MR images were acquired, separated the two fMRI runs.

7.2.3 MR data acquisition

The scanning protocol is shown in [Figure 7.1](#). In addition to task-based fMRI, participants underwent a resting state fMRI scan, a diffusion-weighted scan, a structural scan and two MR spectroscopy acquisitions. The present chapter focuses on the first 4 MR modalities. The MR spectroscopy data are reported elsewhere ([chapter 6](#)). The scanning session was approximately 60 minutes long and was well tolerated by all patients and age-matched healthy individuals. Participants wore earplugs throughout scanning. All MR data were acquired on a 3T Siemens Trio MR scanner with a maximum gradient strength of 40 mT.m⁻¹ at the Oxford University Centre for Clinical Magnetic Resonance Imaging (OCMR). Details of each sequence are provided in the following sub-sections.




Figure 7.1: MRI scanning protocol. All participants ($n = 29$) underwent a 60 minutes scan involving a resting state scan (RSN), two runs of BOLD fMRI acquisition while participants performed the Posner-like task, a T1-weighted structural scan (T1), two MR spectroscopy acquisitions (in M1 and occipital lobe), and a diffusion-weighted scan (DWI).

7.2.3.1 Resting state fMRI

Resting state fMRI data were acquired (180 volumes; 34×3.5 mm slices; $TR/TE = 2000/28$ ms; flip angle = 89 deg; $FOV = 192$; voxel size = $3 \times 3 \times 3.5$ mm³) during 6 minutes at the beginning of each scanning session. Participants were instructed to keep their eyes open and fixate on a cross hair projected on the screen.

7.2.3.2 Posner-like task fMRI

All participants underwent two task fMRI runs with a break in-between corresponding to the time required to acquire the T1-weighted images (approximately 9 min). During each fMRI run, 230 axial echo-planar volumes (34×3.5 mm slices; $TR/TE = 2000/28$ ms; flip angle = 89 deg; $FOV = 192$; voxel size = $3 \times 3 \times 3.5$ mm³) were acquired over 7 min 44 secs while participants performed 8 blocks of a Posner-like task. At the end of the second fMRI run, a GRE field map (44×3.5 mm axial slices, $TR = 488$ ms, $TE1/2 = 5.19/7.65$ ms, flip angle = 60 deg, $FOV = 224$, voxel size = 3.5 mm isotropic, scan duration = 63 secs) was acquired for correction of geometric distortions caused by magnetic field inhomogeneity.

7.2.3.3 T1-weighted structural images

High resolution T1-weighted structural MR images (224×1 mm axial slices; $TR/TE = 3000/4.71$ ms; flip angle = 8 deg; $FOV = 256$; voxel size = 1 mm isotropic; scan time = 528 secs) were acquired after the first task fMRI run for MRS voxel placement,

registration, and lesion masking purposes.

7.2.3.4 Diffusion weighted imaging

Diffusion-weighted echo planar images (1 acquisition of 64 diffusion-weighted directions, $b\text{-value} = 1500 \text{ s/mm}^2$; 5 non-diffusion weighted images with 1 b_0 volume in the reverse phase-encode direction, $b\text{-value} = 0 \text{ s/mm}^2$; 64 slices; $\text{TR/TE} = 8900/94.8 \text{ ms}$; $\text{FOV} = 192 \times 192$; voxel size = 2 mm isotropic; $\text{PAT} = 2$; scan time = 596 secs) were acquired at the end of the scanning session for all participants.

7.2.4 MR data analysis

7.2.4.1 Lesion mapping

Lesion masks, including white matter and grey matter, were delineated manually on the T1-weighted structural images and then registered onto the standard 1 mm Montreal Neurological Institute (MNI) space. To characterise the white matter pathways likely disrupted by patients' lesions, the Disconnectome map software, part of the BCB toolkit, was used (<http://www.brainconnectivitybehaviour.eu>) (Lunven *et al.* 2015; Thiebaut de Schotten *et al.* 2015). The registered lesion masks for each patient were used as seed points to identify which tracts passed through the damaged regions, by comparison with a reference dataset of healthy controls. Each patient's disconnectome map was thresholded at a high value of 90% (i.e. probability of disconnection $\geq 90\%$) and then overlaid in order to identify the tracts that were most likely to be disconnected in all six patients (region in dark grey in Figure 7.2B).

7.2.4.2 Posner-like task fMRI data

All image processing and statistical analyses were performed using tools from the Oxford Centre for Functional Magnetic Resonance Imaging of the Brain Software Library (FSL version 5.0, Oxford, U.K; www.fmrib.ox.ac.uk/fsl). The first three

volumes of each fMRI run were discarded to allow for T1 equilibration effects. The fMRI data were then corrected for motion using MCFLIRT (Jenkinson *et al.* 2002) and non-brain structures were removed using BET (Smith 2002). The images were spatially smoothed using a 6 mm Gaussian kernel of full-width at half maximum and low-frequency drifts were removed with high-pass temporal filtering with a cut-off period of 60 seconds. Geometrical distortions induced by magnetic field inhomogeneities were corrected using the field map image with FUGUE (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FUGUE>). Functional data were registered to individuals' structural scan using boundary-based linear registration with a de-weighting of voxels included in the lesion mask (Jenkinson & Smith 2001; Jenkinson *et al.* 2002; Greve & Fischl 2009). Next, individuals' structural scan was registered to the 2 mm MNI152 standard space using non-linear registration (Andersson *et al.* 2007) with a de-weighting of voxels included in the lesion mask. Registration of functional data to the standard space was obtained by combining the linear and non-linear transformations.

FEAT (version 6.0) was used to perform a general linear model analysis on a voxelwise basis. Time series statistical analysis was carried out using FILM with local auto-correlation correction (Woolrich *et al.* 2001). A 2-explanatory variables (EVs) design was fitted to the BOLD time series at the individual level. Task blocks were modelled as one EV to capture the main effect of interest, i.e. the deployment of visuospatial attention during the Posner-like task. Additionally, button presses were modelled as a stick function to capture extra variance introduced by the motor response. The haemodynamic response was modelled as a gamma function, a normalisation of the probability density function of the gamma distribution with zero phase, standard deviation of 3 seconds, and a mean lag of 6 seconds.

To compare the activation across the neglect and control groups, a mixed-effects

analysis was conducted. At the first level, Z statistic images were generated for the task contrast, separately for the first and second run. Then, a second-level fixed-effects analysis was conducted on each participant's pair of scans to yield a Z statistic map reflecting the average pattern of brain activity across the two runs. Finally, a third level mixed-effects analysis (FLAME 1+2, with automatic outlier de-weighting) was conducted at the group level on the output from the second level analysis, with Z statistic images thresholded using clusters determined by $Z \geq 3.1$ (corresponding to $p \leq 10^{-3}$), and a corrected cluster significance threshold of $p = 0.05$. The activation maps of each group were considered separately as well as contrasted against one another.

7.2.4.3 Resting state fMRI data

The resting state analysis was focused on identifying and comparing the frontoparietal networks, known to be relevant for visuospatial attention (Smith *et al.* 2009; Markett *et al.* 2014), between neglect patients and age-matched healthy controls. Preprocessing of resting state data involved the same steps as the Posner-like fMRI data (first 3 volumes removed, brain extraction, 6 mm spatial smoothing, 100 secs high-pass filtering, B0 unwarping, registration to MNI template with de-weighting of the lesion). Because of unequal sample sizes ($n = 23$ age-matched healthy individuals; $n = 6$ neglect patients), defining the resting state networks (RSNs) based on a group independent component analysis (ICA) of the 29 individuals (age-matched healthy and patients as a single group) would bias estimates in favour of the largest group (i.e. age-matched healthy controls). This would introduce a systematic confound when comparing the two groups. To avoid this issue, we instead used the 13 RSNs estimated in a separate group in [chapter 5](#) ($n = 22$, 0 female, mean age = 24.9, SD = 3.3). This approach also ensured direct comparability with the RSNs identified as relevant for long-term retention of prism after-effects ([chapter 5](#)). This set of spatial maps was used to

generate subject-specific versions of the spatial maps, and associated timeseries, using dual regression (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/DualRegression>) (Beckmann *et al.* 2009a; Filippini *et al.* 2009). The analysis was as follows. First, for each participant, the group-average set of spatial maps was regressed (as spatial regressors in a multiple regression) into the subject’s 4D space-time dataset. This resulted in a set of subject-specific timeseries, one per group-level spatial map (i.e. 13 per subject). Next, those timeseries were regressed (as temporal regressors, again in a multiple regression) into the same 4D dataset, resulting in a set of subject-specific spatial maps, one per group-level spatial map (i.e. 13 per subject). We then tested for group differences using accelerated permutation inference implemented in FSL’s Permutation Analysis of Linear Models tool (threshold-free cluster enhancement, family-wise error correction) and correcting across contrasts and modalities (PALM version alpha-109; <https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/PALM>) (Winkler *et al.* 2014; Winkler *et al.* 2016a; Winkler *et al.* 2016b).

The presence of a lesion in one group raises interpretative challenges because of systematic, non-specific differences it may induce in RSNs that most strongly overlap with the lesion (e.g. right frontoparietal network). The test of inter-group differences therefore focused on the left frontoparietal network only, because it was the only RSN predominantly located in the healthy, non lesioned hemisphere in neglect patients.

7.2.4.4 Diffusion-weighted data

Diffusion-weighted images were processed using tools from the FMRIB Software Toolbox (FSL; <https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/>). Susceptibility-induced, subject motion and eddy current distortions were corrected using TOPUP (Andersson *et al.* 2003; Smith *et al.* 2004) and EDDY (Andersson & Sotiropoulos 2016). Tractography was run using ProbTrackx (Behrens *et al.* 2003; Behrens *et al.* 2007a)

to explore the pattern of structural connectivity of certain seeds of interest, identified as relevant to neglect patho-physiology through the analysis of the Posner-like fMRI data. Tractography was run in individuals' native space using 5000 samples, then registered to the MNI standard space using the same transformation as for the functional data. Group mean path distribution estimate maps were thresholded at a high value of 1000 for visualisation.

7.3 Results

7.3.1 Lesion mapping

All patients had suffered a right hemispheric stroke and had an intact left hemisphere. [Table 7.1](#) lists all infarcted regions for each patient and [Figure 7.2A](#) shows the lesion overlap map in the standard space. The only region of overlap between all 6 patients was located in the right putamen and surrounding white matter (region in yellow in [Figure 7.2A](#)).



A) Lesion overlap map**B) Disconnectome map**

Figure 7.2: Patients' lesion anatomy. **A.** Lesion overlap between the six patients in standard MNI space. A precise description of the lesion of each patient is provided in [Table 7.1](#). Colour bar represents the number of overlapping lesions. **B.** 'Disconnectome map' shows the white matter tracts with high probability of disconnection (>90%) in all six patients (in yellow). They include the right inferior fronto-occipital fasciculus (IFOF, in red), the right superior longitudinal fasciculus (SLF, in green), the right anterior thalamic radiations (ATR, in blue) and the right corticospinal tract (CST, in black). The tract masks are taken from the JHU DTI-based white matter tractography atlas (Mori *et al.* 2005). R: right hemisphere. All images are displayed in radiological convention (i.e. left side of the image corresponds to the right side of the brain). MNI coordinates are displayed under each slice.

As outlined earlier in this chapter (section 7.1), the current dominant hodotopic patho-physiological model of spatial neglect insists on the notion that local lesions have network level effects (Doricchi *et al.* 2008). In order to characterise the consequence of the lesion on white matter pathways, a 'disconnectome' map was generated for each patient using the BCB toolkit (Thiebaut de Schotten *et al.* 2015). In brief, this analysis involves identifying white matter tracts that are likely interrupted by the lesion based on the spatial distribution of the lesion (delineated on the T1-weighted structural scan).

Lesioned area	P1	P2	P3	P4	P5	P6
Precentral gyrus	✓			✓	✓	
Superior frontal gyrus					✓	
Middle frontal gyrus	✓	✓		✓	✓	
Inferior frontal gyrus	✓	✓		✓	✓	
Medial frontal gyrus					✓	
Insula	✓		✓	✓	✓	✓
Postcentral gyrus	✓			✓	✓	
Superior parietal lobule	✓	✓			✓	
Inferior parietal lobule	✓	✓	✓	✓	✓	
Precuneus	✓					
Superior temporal gyrus	✓	✓	✓	✓	✓	
Middle temporal gyrus	✓	✓	✓	✓	✓	
Inferior temporal gyrus		✓				
Temporal pole		✓	✓	✓		
Occipitotemporal gyrus		✓				
Parahippocampal gyrus		✓				
Superior occipital gyrus		✓			✓	
Middle occipital gyrus	✓	✓	✓	✓		
Inferior occipital gyrus		✓				
Occipital pole		✓				
Cuneus		✓				
Lingual gyrus		✓				
Fusiform gyrus		✓				
Putamen	✓	✓	✓	✓	✓	✓
Globus pallidus				✓	✓	
Thalamus				✓	✓	

Table 7.1: Patients' lesion anatomy

This approach revealed a greater degree of overlap between all 6 patients in terms of the consequences of the lesion on structural connectivity ([Figure 7.2B](#)). Four tracts were likely interrupted in the right hemisphere in all 6 patients: the anterior thalamic radiations (ATR), the corticospinal tract (CST), the IFOF and the SLF (predominantly branches I & II). The last two tracts have previously been associated to spatial neglect (Thiebaut de Schotten *et al.* 2005; Roux *et al.* 2011; Urbanski *et al.* 2011; Bozzali *et al.* 2012; Thiebaut de Schotten *et al.* 2012; Vallar *et al.* 2014; Lunven *et al.*

2015; Herbet *et al.* 2017). They connect key nodes of the cortical network involved in the control of visuospatial attention within right occipital, parietal and frontal regions (Thiebaut De Schotten *et al.* 2011; Forkel *et al.* 2014).

7.3.2 Re-organisation of the left frontoparietal network in spatial neglect

In order to characterise the effect of the lesion on the left frontoparietal attention network (FPN), situated predominantly in the intact hemisphere, a dual regression analysis was conducted.




Figure 7.3: Functional re-organisation of the left frontoparietal network. **A.** Left frontoparietal network comparison between neglect patients and age-matched controls revealed increased co-activation in bilateral fusiform gyrus (FG), left lateral occipital cortex (LOC), left posterior middle temporal gyrus (MTG), and middle section of the left superior temporal sulcus (STS). Red-to-yellow colours define increases in co-activation ($p < 0.05$, threshold-free cluster enhanced, family-wise error corrected). **B.** The same analysis revealed a sub-threshold decreased co-activation with the left FPN (threshold-free cluster enhanced, not family-wise error corrected, $p < 0.001$) within right IPS/IPL at the junction between angular gyrus and supramarginal gyrus. Colour code represent uncorrected $-\log(p)$. R: right hemisphere. All images are displayed in radiological convention (i.e. left side of the image corresponds to the right side of the brain). MNI coordinates are displayed under each slice.

Compared to age-matched healthy controls, neglect patients showed a significantly

increased degree of co-activation with the left FPN (measured by the regression coefficient as part of the dual regression approach) in the bilateral fusiform gyrus, left lateral occipital cortex, left middle temporal gyrus and left superior temporal sulcus (Figure 7.3A). This indicates that the left FPN was functionally more connected to visual areas bilaterally and regions of the VAN in the left hemisphere in neglect, potentially indicating compensatory mechanisms. After correction for multiple comparison, no brain area showed a significant decreased level of co-activation with the left FPN in neglect. However, inspection of the uncorrected map with a high significance threshold (threshold-free cluster enhanced, not family-wise error corrected, $p < 0.001$; Figure 7.3B) revealed a decrease in the degree of co-activation with the left FPN in the right IPS/IPL at the junction between angular gyrus and supramarginal gyrus, in a region labelled as cluster D in Mars and colleagues' tractography-based parcellation (Mars *et al.* 2011b). (Note that this parietal cluster largely overlaps with the ANG/IPS cluster within which inter-hemispheric dominance was modulated by right inhibitory rTMS in Appendix B (Figure B.4).

Unfortunately, the anatomical specificity of this result could not be evaluated because the left FPN was the only network lateralised to the non-lesioned hemisphere. All the other RSNs were predominantly bilateral and overlapped with the lesion to a certain extent in stroke patients. This prevented us from using any of them as a 'control' network, leaving open the question of the specificity of the effect aforementioned.

7.3.3 Behavioural performance on the Posner-like task

Functional MRI data were collected while participants performed a Posner-like visual detection task inside the scanner. Behavioural performance was assessed by hit rates (HR), mean reaction time (RT) and false alarm rate (FAR) (Figure 7.4). Because of

unequal sample sizes, statistical differences between age-matched healthy individuals and neglect patients were analysed using non-parametric Mann–Whitney–Wilcoxon testing (Mann & Whitney 1947; Fay & Proschan 2010; MacFarland & Yates 2016), run separately in the left and right visual hemifield, with a Bonferroni significance threshold of $p < 0.008$. Both groups detected visual targets reliably in the right hemifield and performance did not differ significantly between patients and controls (hit rate $> 80\%$ for all individuals; $Z = 0.17$, $p = 0.84$; Figure 7.4). However, accuracy in the left hemifield was impaired in patients (patients: mean HR = 62.0%, SEM = 14.5%, age-matched controls: mean HR = 97.9%, SEM = 0.9%; $Z = 3.68$, $p < 10^{-3}$ Figure 7.4A). To quantify spatial bias in performance, i.e. the core defining clinical feature of spatial neglect (Heilman & Van Den Abell 1980), a lateralisation index (LI) was computed as follows:

$$LI = \frac{(Right - Left)}{(Right + Left)} \quad (7.1)$$

where 'Right' and 'Left' refer to the HR in the respective hemifield. As expected, neglect patients' accuracy was significantly more right-biased than that of healthy controls (patients: mean LI = 0.28, SEM = 0.14, age-matched controls: mean LI = $9.8 \cdot 10^{-4}$, SEM = $3.2 \cdot 10^{-3}$; $Z = 3.60$, $p < 10^{-3}$; Figure 7.4A).




Figure 7.4: Behavioural performance inside the scanner. **A.** Left graph shows percentage correct reports (hit rate) across the left and right visual hemifields for neglect patients (in dark grey) and age-matched healthy individuals (in white) (group mean, error bar = SEM). There was no difference in accuracy in the right hemifield ($p = 0.84$). Neglect patients' performance was impaired in the left hemifield ($p < 10^{-3}$). Right graph shows group mean lateralisation index (error bar = SEM) calculated on hit rate (i.e. positive values mean more targets detected in the right than left hemifield). Neglect patients were significantly more right biased than age-matched healthy individuals ($p < 10^{-3}$). **B.** Left graph shows reaction time (in ms) across the left and right visual hemifields for neglect patients (in dark grey) and age-matched healthy individuals (in white) (group mean, error bar = SEM). Overall, patients were slower across both hemifields (both $p < 0.003$). Right graph shows group mean lateralisation index (error bar = SEM) calculated on reaction time (i.e. negative values indicate longer reaction times for targets on the left than for those on the right). Patients' RT lateralisation patterns did not significantly differ from those of age-matched healthy individuals ($p = 0.25$).

Patients had significantly slower RTs compared to age-matched controls, across both visual hemifields (left: $Z = 3.04$, $p = 0.001$; right: $Z = 2.77$, $p = 0.003$; Figure 7.4B). The difference in RT between left and right hemifields – as assessed by a LI computed on RT – did not differ between patients and controls (patients: mean

LI = -0.08, SEM = 0.05, age-matched controls: mean LI = -0.01, SEM = 0.005; $Z = 1.16$, $p = 0.25$; [Figure 7.4B](#)). The false alarm rate was low ($< 12.5\%$ in all individuals) and did not differ between the two groups (patients: mean FAR = 3.65, SEM = 2.16, age-matched controls: mean FAR = 0.61, SEM = 0.16; $Z = 0.99$, $p = 0.32$).

7.3.4 Altered inter-hemispheric balance in spatial neglect

To identify the network of brain regions activated during the deployment of visuospatial attention, a mixed-effects analysis was conducted separately for age-matched healthy individuals and neglect patients. The pattern of functional activation during the Posner-task is first described in age-matched healthy controls ([section 7.3.4.1](#)) and neglect patients ([section 7.3.4.2](#)) separately, before being contrasted against one another ([section 7.3.4.3](#)).

7.3.4.1 Age-matched healthy individuals

Whole-brain voxel-wise task-related activation maps (motor-related activity regressed out; $Z \geq 3.1$, corrected cluster significance threshold of $p = 0.05$) are shown for age-matched healthy individuals in [Figure 7.5A](#). Anatomical localisation of active areas was determined using the Harvard-Oxford cortical and subcortical atlases ([Frazier et al. 2005](#); [Desikan et al. 2006](#)) as well as connectivity-based cortical parcellations ([Beckmann et al. 2009b](#); [Mars et al. 2011a](#); [Mars et al. 2011b](#); [Sallet et al. 2013](#)).

During the Posner-like task, age-matched healthy individuals revealed widespread activation within the visual system, and ventral and dorsal attention networks ([Corbetta et al. 1998](#); [Corbetta & Shulman 2002](#); [Corbetta & Shulman 2011](#); [Umarova et al. 2011](#)). Activation within the visual system included fusiform gyrus (FG) and occipital pole. Activation within the DAN included bilateral lateral occipital cortex (LOC), intraparietal sulci and superior parietal lobules (IPS/SPL), supramarginal

gyrus (SMG) and frontal eye field (FEF). (Note that in age-matched healthy volunteers, the activation in the IPS overlaps with the ANG/IPS cluster within which inter-hemispheric dominance was modulated by right inhibitory rTMS in [Appendix B](#) ([Figure B.4](#)). Activation within the VAN included temporoparietal junction (TPJ; posterior division of the angular gyrus and middle temporal gyrus), pre-central gyrus (PCG), inferior frontal gyrus (IFG) and anterior insula (Ins).




Figure 7.5: Activation during the Posner-like task. Regions activated during the Posner-like task (button-press related activity regressed out) in age-matched healthy individuals ($n = 23$, **A**) and chronic spatial neglect patients ($n = 6$, **B**). Cerebellar activation is represented in slices rather than 3D rendering. Voxelwise statistics are shown for $Z \geq 3.1$ with a corrected cluster significance threshold of $p = 0.05$. Colour bar represents Z statistics at each voxel. **B.** Group lesion overlap is shown in dark grey in the right hemisphere. ACC: Anterior cingulate cortex; ANG: Angular gyrus/Intraparietal sulcus; FEF: Frontal eye field; FG: Fusiform gyrus; FP: Frontal pole; IFG: Inferior frontal gyrus; Ins: Insula; IPS/SPL: Intraparietal sulcus/Superior parietal lobule; LOC: Lateral occipital cortex; M1: Primary motor cortex; MTG: Middle temporal gyrus; Postcentral: Postcentral gyrus; Precentral: Precentral gyrus; SMA: Supplementary Motor Area; VI: cerebellar lobule VI. R: right hemisphere. All slices are displayed in radiological convention (i.e. left side of the image corresponds to the right side of the brain). MNI coordinates are displayed under each cerebellar slice.

Finally, bilateral activation was also detected in the superior posterior cerebellum (Crus 1 and lobule VI, [Figure 7.5A](#)). This activation is unlikely to reflect motor-related activity, which was detected more anteriorly ([Figure 7.6](#)).




Figure 7.6: Motor-related activation in age-matched healthy individuals. **A.** Regions showing button press related activation during the Posner-like task (task block activity regressed out) are shown for age-matched healthy individuals. Voxelwise statistics are shown for $Z \geq 3.1$ with a corrected cluster significance threshold of $p = 0.05$. Colour bar represents Z statistics at each voxel. **B.** For comparison, the 'leg' (in yellow), 'face' (in blue) and 'hand' (in green) sensorimotor RSN maps from [chapter 5](#) are shown ($Z \geq 4$). R: right hemisphere. All slices are displayed in radiological convention (i.e. left side of the image corresponds to the right side of the brain). MNI coordinates are displayed under each cerebellar slice.

As outlined in [Appendix B](#), lateralisation of activity within the visuospatial attention network constitutes an important signal for the allocation of attention in space (Kinsbourne 1976; Duecker & Sack 2015; Petit et al. 2015). Coherent with the current view that the ventral attention network is predominantly localised to the right hemisphere (Corbetta & Shulman 2002; Corbetta et al. 2008; Corbetta & Shulman 2011), observation of thresholded activation maps ([Figure 7.5A](#)) suggested a certain degree of functional lateralisation within the right TPJ and inferior frontal regions. However, this apparent right dominance could be an artefact resulting from the thresholding of functional maps (Macaluso & Patria 2007; Shulman et al. 2010; Agcaoglu et al. 2015). In order to overcome this methodological issue and formally assess inter-hemispheric dominance, we used the same analysis method as described in [Appendix B](#), i.e. we computed the difference in activity between homologous voxels of the two hemispheres and tested for significant deviations from zero using permutation

testing (Anderson & Robinson 2001; Winkler *et al.* 2014). The resulting voxelwise functional lateralisation maps (threshold-free cluster enhanced, family-wise error corrected, $p < 0.05$) are shown for age-matched healthy individuals in Figure 7.7A.




Figure 7.7: Functional lateralisation during the Posner-like task. A. Regions showing significant hemispheric lateralisation of activation during the Posner-like task in age-matched healthy elderly individuals ($n = 23$, **A**) and chronic spatial neglect patients ($n = 6$, **B**). Voxelwise statistics of functional lateralisation are shown for family-wise error (FWE) corrected $p < 0.05$ with threshold-free cluster enhancement. Colour bar represents FWE-corrected p -value at each voxel. Right functional lateralisation is shown in the right hemisphere and left functional lateralisation in the left hemisphere. **A.** In age-matched healthy individuals, activation during the Posner-like task was found to be left lateralised within the primary motor cortex (M1) and right lateralised within the posterior subdivision of the temporoparietal junction (TPJp), posterior section of the middle frontal gyrus, and anterior cingulate gyrus (ACC). **B.** In chronic spatial neglect patients, owing to the lesion, no significant right functional lateralisation was found. Instead, significant left functional lateralisation was detected in the lateral occipital cortex (LOC).

This analysis confirmed the right lateralisation of some nodes of the ventral attention network. Right hemisphere dominance was indeed detected in the posterior subdivision of the TPJ (Mars *et al.* 2011a), in a region previously identified as showing right hemisphere dominance both during stimulus-driven shifts of spatial attention

and during target detection at an attended location (Shulman *et al.* 2010). Two other regions showed significant right hemisphere dominance: the posterior section of the middle frontal gyrus (cluster 8 in Sallet and colleagues' parcellation (Sallet *et al.* 2013)) and the middle section of the cingulate gyrus (cluster 4 in Beckmann and colleagues' parcellation (Beckmann *et al.* 2009b)). Note that none of the DAN nodes (LOC, IPS, FEF) showed functional lateralisation in healthy controls, indicating balanced inter-hemispheric competition in this functional network.

Finally, left hemisphere dominance was detected in the primary motor cortex (M1), likely reflecting the fact that participants used their right hand to report the presence of targets throughout the task (Yousry *et al.* 1997).

7.3.4.2 Neglect patients

In neglect patients, the overall pattern of brain activation was comparable, except from apparent lack of activation within the lesioned right hemisphere (Figure 7.5B). The Posner-like task activated the right occipital pole, bilateral fusiform gyrus, bilateral LOC, left IPS/SPL, left SMG, left anterior insula, bilateral FEF, left M1 and left postcentral gyrus. Unlike age-matched healthy individuals, neglect patients showed no significant right dominance in their pattern of activation during the Posner-like task. Instead, significant left dominance was found in the LOC (Figure 7.7B).



7.3.4.3 Neglect patients versus age-matched healthy controls

Whole brain two-sample unpaired *t*-test ($Z \geq 3.1$, corrected cluster significance threshold of $p = 0.05$) was used to compare task-related BOLD signal changes between the two groups.

Area	Side	MNI coordinates			Z-value
		<i>x</i>	<i>y</i>	<i>z</i>	
<i>Controls > Patients</i>					
Precentral gyrus	L	-38	-8	42	5.89
Superior posterior cerebellum	L	-10	-78	-26	5.01
<i>Patients > Controls</i>					
Lateral occipital cortex	L	-42	-88	8	6.77
Superior parietal lobule	L	-54	-36	54	6.06
Postcentral gyrus					
Fusiform gyrus	L	-32	-72	-8	5.60

Table 7.2: Brain regions showing differences in activity between neglect patients and age-matched healthy individuals. Coordinates are the peak voxel for the two-sample unpaired t-test contrasting the task-related activity (motor-related activity regressed out) across the two groups.

Compared to age-matched healthy controls, patients with chronic neglect demonstrated an increased level of activation within the left LOC, left fusiform gyrus, left postcentral gyrus, and left SPL/IPS (regions in red in [Figure 7.8A](#), [Table 7.2](#)), while no increase in activity was detected in the right hemisphere. Decreased activity relative to age-matched controls was identified in neglect patients in the left precentral gyrus and left superior posterior cerebellum (Crus 1 and lobule VI) (regions in blue in [Figure 7.8A](#), [Table 7.2](#)). (Note that the precentral cluster was located underneath where the anode electrode was positioned during prism adaptation in the next chapter ([chapter 8](#))).


A) Change in activation map**B) Change in lateralisation map****Figure 7.8: Comparison of brain activity during the Posner-like task.**

A. Regions showing greater activation in spatial neglect patients compared to age-matched healthy controls during the Posner-like task are shown in red. Regions showing lower activation in patients compared to controls are shown in blue. Voxelwise statistics (two-sample unpaired t -test) are shown for $Z \geq 3.1$ with a corrected cluster significance threshold of $p = 0.05$. Colour bar represents Z statistics at each voxel. **B.** Regions showing greater functional lateralisation during the Posner-like task in neglect patients compared to age-matched healthy individuals. Map shows the only region (LOC) in which there was greater left hemispheric dominance in neglect patients relative to controls (FWE-corrected $p < 0.05$, threshold-free cluster enhancement). No other difference in lateralisation (e.g. greater right hemisphere dominance in neglect patients) was found. Colour bar represents FWE-corrected p -value at each voxel. Tractography from this region is shown in Figure 7.9. All slices are displayed in radiological convention (i.e. left side of the image corresponds to the right side of the brain). MNI coordinates are displayed under each cerebellar slice. R: right hemisphere. FG = fusiform gyrus, LOC = lateral occipital, Postcentral = Postcentral gyrus, Precentral = Precentral gyrus, psCer = posterior superior cerebellum, SPL/IPS = superior parietal lobule/intraparietal sulcus.

To evaluate whether these unilateral differences in activity resulted in significant alterations of the inter-hemispheric balance in neglect patients, the functional lat-

eralisation maps of the two groups were compared using a whole brain two-sample unpaired t -test implemented in randomise (family-wise error correction, threshold free cluster enhancement) (Anderson & Robinson 2001). This analysis revealed an increase in left dominance within the LOC (Figure 7.8B), in a cluster that was at least partially spared in all patients except patient 2.

A) Healthy individuals



B) Stroke patients




Figure 7.9: Results of a tractography run from the region showing increased left dominance in neglect patients. Tractography (5000 samples) was run from the homologous LOCs found to show exaggerated left dominance in neglect patients (Figure 7.8B). Group mean path distribution estimate maps (threshold = 1000) are shown for age-matched healthy individuals (A) and neglect patients (B). R: right hemisphere. All slices are displayed in radiological convention (i.e. left side of the image corresponds to the right side of the brain). MNI coordinates are displayed under each slice.

Finally, in order to explore white matter disconnections that may underlie this abnormal left dominance within LOC, we used this cluster and its homologous contralateral right cluster as seeds to run tractography in patients and age-matched healthy individuals. The results are shown in Figure 7.9 (group mean path distribu-

tion estimate threshold = 1000). In age-matched healthy individuals, this revealed a widespread bilateral white matter network including the forceps major, the IFOFs, the SLFs and the corticospinal tracts (Figure 7.9A). LOCs therefore appeared to be structurally connected to 1) their homologous counterpart; 2) other cortical nodes of the DAN, including IPS/SPL and FEF; and 3) cortical nodes of the VAN including the anterior insula and middle temporal gyri. The connectivity pattern in the non-lesioned left hemisphere was comparable in patients. In all patients except patient 2, the communication between the two LOC clusters, via the forceps major, was intact. However, in the right hemisphere, *in vivo* tractography dissection of the connections of this cluster revealed a profound disconnection with more anterior regions due to the presence of the lesion (Figure 7.9B).

7.4 Discussion

This study aimed to characterise the neural basis of post-stroke chronic spatial neglect in a sample of six patients who participated in the rehabilitation study described in chapter 8. Building on previous findings that bilateral visuospatial attentional depends upon the balance of functional activity between left and right parietal cortices in the healthy brain (Appendix B; Petitet *et al.* 2015), we expected chronic spatial neglect to be associated with a left hemispheric dominance within regions of the dorsal attention network (Corbetta & Shulman 2002; Corbetta & Shulman 2011). To test this, participants ($n = 6$ neglect patients, $n = 23$ age-matched healthy individuals) performed a target detection task (Posner-like task) while BOLD fMRI data was acquired. The scanning session also included a T1-weighted structural scan, a resting state fMRI scan, a diffusion-weighted scan and MR spectroscopy. The lesion distribution in the patient group was consistent with the literature. White matter tracts interconnecting occipital, parietal and frontal regions within the right hemi-

sphere (SLF, IFOF) were found to be likely interrupted in all patients (Figure 7.2 & 7.9). The consequences of the lesion on the resting state networks included increased functional coupling of the frontoparietal attention network in the intact left hemisphere with bilateral visual areas (Figure 7.3). In age-matched healthy individuals, the Posner-like task activated a widespread bilateral visuospatial attention network (Umarova *et al.* 2011) that was predominantly right-lateralised within regions of the VAN (TPJ, precentral gyrus) (Corbetta & Shulman 2011). No functional lateralisation however was detected within regions of the DAN (Figure 7.7), consistent with non-biased behavioural performance of healthy individuals in the scanner (Figure 7.4). By contrast to age-matched healthy individuals, neglect patients showed greater activation within the left postcentral gyrus, SPL/IPS, LOC and fusiform gyrus (Figure 7.8A). Their behavioural performance inside the scanner was also significantly more biased towards the right ipsilesional hemifield. Voxelwise analysis of functional lateralisation revealed that neglect was associated with greater left dominance within the LOC, a key-node of the dorsal attention network (Kinsbourne 1976; Cohen *et al.* 1994).

The acquisition and analysis of neuroimaging data in stroke patients are challenging for several reasons. In order to maximise data quality and participants' comfort, the overall scanning time was approximately 1 hour only (Figure 7.1). The resting-state and task fMRI scans were acquired early during the session in order to avoid confound by fatigue. In addition, a block design was used for task fMRI (Umarova *et al.* 2011), enabling us to measure the overall brain activity during the Posner-like task, but preventing us from characterising further the role of each region in more precise psychological terms (which would require an event-related design, see chapter 3). Another methodological limitation of this study is the absence of eye-tracking data inside the scanner. All participants were trained outside the scanner to maintain

fixation on the cross hair at all time but visual fixation could not be monitored inside the scanner for technical reasons. Thus, it is possible that gaze-shifts to the ipsilesional right visual hemifield, one of the clinical features of neglect, may contribute to differences in brain activity between the two groups. However, this type of spatial bias in eye movements has predominantly been reported in visual observation and visual search tasks, in which patients are passively or actively exploring a visual scene (Karnath 1994; Karnath & Fetter 1995; Husain *et al.* 2001; Duc *et al.* 2008; Bays *et al.* 2010). By contrast, when performing a cued visual detection task outside the scanner (Posner cueing task in [chapter 8](#)), the neglect patients included in this study did not show any systematic fixation failure.

The cerebellum has a long history as a neural structure viewed to be implicated exclusively in motor control. However, the last three decades have provided some evidence of its involvement in non-motor cognitive functions such as emotional processing, associative learning or attention (Allen & Courchesne 1998; C. Justus & Ivry 2001; Stoodley & Schmahmann 2009; Lange *et al.* 2015; Strata 2015). Although several neuroimaging studies reported non-motor related cerebellar activation within the left posterior superior cerebellum during visuospatial and attention tasks in the healthy brain (Allen *et al.* 1997; Le *et al.* 1998; Fink *et al.* 2000; O'Reilly *et al.* 2008), the lesion of this region does not typically lead to strong attentional deficits (Akshoomoff & Courchesne 1994; Yamaguchi *et al.* 1998; Ravizza & Ivry 2001, however, see Silveri *et al.* 2001). Thus, the contribution of this region to neglect has, to date, remained unclear. To our knowledge, our data are the first to provide direct evidence of a pathological hypo-activation within the left superior posterior cerebellum (Crus 1, lobule VI) in neglect patients ([Figure 7.8](#)). This hypo-activation appears to be an indirect effect of the lesion, caused by functional reorganisation in a region distant from the lesion. This functional impairment is unlikely to reflect a motor deficit for

the following reasons. First, in our analysis, the motor-related activity was regressed out from the overall task activity by including an explanatory variable modelling button presses to the general linear model design matrix. Inspection of the button press related activation map ([Figure 7.6](#)) revealed a sensorimotor network including predominantly the left M1/S1 and the right anterior cerebellum (network reassembling the 'hand' sensorimotor resting state network depicted in green in [Figure 7.6B](#)). By contrast, the anatomical localisation of the cerebellar cluster showing pathological hypo-activation in neglect was more posterior and was located in the left cerebellar hemisphere. Second, this cluster overlapped with the cerebellar node of the right frontoparietal network ([chapter 5](#)), indicating that it belongs to a functional network known to be involved in attention in healthy individuals (Fox *et al.* 2006; Smith *et al.* 2009; Parks & Madden 2013) and to be disrupted by the lesion in neglect patients (He *et al.* 2007). Our previous work ([chapter 5](#)) demonstrated a role of the functional interaction between the right FPN and the sensorimotor network in the generalisation and retention of prism after-effects. Because of its involvement in both neglect pathophysiology and prism after-effects, the left superior posterior cerebellum may play a role in mediating the clinical effect of prism adaptation in neglect. We speculate that it may influence higher order cortical regions of the right frontoparietal network via a bottom-up influence during prism adaptation (Luauté *et al.* 2006).

In our data, the left primary motor cortex (M1) showed the same pattern of decreased activation compared to age-matched controls during the Posner-like task in neglect patients ([Figure 7.8](#)). This is of particular interest in the context of the M1 excitatory stimulation (anodal tDCS) intervention tested in the following chapter ([chapter 8](#)). However, this result should be taken carefully because of the large degree of overlap with the button press related activation ([Figure 7.6](#)) and the large white matter composition of this cluster. Further work should better characterise the

non-motor contribution of the cerebellum and primary motor cortex to visuospatial attention deficits in neglect potentially using tasks that do not require a right hand motor response.

In the present study, the left SPL/IPS, LOC and fusiform gyrus, which were spared and remote from the lesion in all six patients, had a pathologically increased level of activity during the Posner-like task in neglect patients (Figure 7.8A). Consistent with previous findings in the literature (Vuilleumier *et al.* 2001; Corbetta *et al.* 2005; Vuilleumier *et al.* 2008; Umarova *et al.* 2011), this result suggested an imbalance in activity within regions of the visual system and dorsal attention network. In the healthy brain, the top down modulation of early visual areas from the dorsal frontoparietal network is involved in conscious perception (Lumer *et al.* 1998; Bressler *et al.* 2008; Corbetta & Shulman 2011). Typically, spatially selective modulations of activity are observed within the DAN during shifts of visuospatial attention, demonstrating its role in the coding of stimulus saliency (Treue 1996; Treue & Maunsell 1999; Shulman *et al.* 2009; Shulman *et al.* 2010; Corbetta & Shulman 2011; Tosoni *et al.* 2013). The location of the focus of attention is therefore thought to depend causally on the balance of activity within regions of the DAN (Corbetta & Shulman 2011). To formally assess the impact of abnormal left hyperactivity on the inter-hemispheric balance, we compared the degree of functional lateralisation between homologous voxels across the two groups. This analysis identified pathological left dominance within the LOC as being the most consistent substrate of spatial neglect in our sample (Figure 7.8B). This result confirmed that LOC is a necessary node in visuospatial processing (Corbetta & Shulman 2011) and suggested that the disruption of the inter-hemispheric balance (Kinsbourne 1976; Cohen *et al.* 1994) within this region contributes to chronic spatial neglect.

Further studies should investigate the patho-physiological mechanisms responsible for this functional imbalance in regions that are typically not affected by the lesion in neglect patients. Our tractography results suggest that although the white matter pathway interconnecting homologous left and right LOCs (forceps major) is intact, the connection of the right LOC to the right frontoparietal cortex, via SLF II and IFOF (Thiebaut de Schotten *et al.* 2011; Parlatini *et al.* 2017), appears to be interrupted by the lesion (Figure 7.2, Figure 7.9). Also, the left LOC showed enhanced functional connectivity with the left frontoparietal network in neglect patients compared to age-matched healthy controls (Figure 7.3). We propose that, by depriving the right LOC from excitatory inputs originating from distant, higher-order frontal and parietal areas of the right hemisphere, this structural disconnection may contribute to locally reducing the level of activity within the right LOC (Umarova *et al.* 2011), thereby biasing the inter-hemispheric competition towards the left hemisphere.

7.5 Conclusions

The present study investigated the neural substrates of spatial neglect in a sample of six patients who participated in the rehabilitation study described in chapter 8. To our knowledge, it is the first study to investigate inter-hemispheric competition (Kinsbourne 1976; Cohen *et al.* 1994) during a visuospatial attention task on a voxelwise basis in neglect. Pathological left dominance was detected in the lateral occipital cortex in neglect patients, confirming the role of this brain region in the pathophysiology of neglect (Umarova *et al.* 2011). In addition, neglect was associated with pathological hypo-activation in a region of the left cerebellum known to be involved in non-motor processes (Allen *et al.* 1997; Le *et al.* 1998; Fink *et al.* 2000; O'Reilly *et al.* 2008). Hence, by highlighting the direct involvement of the cerebellum in the pathophysiology of neglect, our data support the hypothesis that modulation of cerebellar

activity, via prism adaptation, may play a role in neglect recovery (Luauté *et al.* 2006; Striemer & Danckert 2010).

Chapter 8

Combining prism adaptation with anodal transcranial direct current stimulation to improve neglect rehabilitation

This final experimental chapter builds on findings from Chapters 5 & 6 and tests the hypothesis that standard prism adaptation therapy combined with sensorimotor anodal transcranial direct current stimulation would enhance the magnitude and/or duration of clinical improvements in chronic neglect patients.

8.1 Introduction

Stroke is a worldwide leading cause of complex disability in adults (Wolfe 2000; Adamson *et al.* 2004; Feigin *et al.* 2014). At the acute stage, a large proportion of right hemisphere stroke survivors suffer from a complex syndrome commonly termed unilateral spatial (hemi-)neglect or (hemi-)inattention (Stone *et al.* 1993; Karnath *et al.*

2004; Milner & McIntosh 2005; Verdon *et al.* 2009). The core defining clinical feature of this syndrome is the patient's inability to attend to, respond adequately to, or orient voluntarily to stimuli located in the contralesional left side of space (Heilman & Van Den Abell 1980; Mesulam 1981; Heilman *et al.* 2002). It refers to a variety of clinical symptoms that are often challenging to assess clinically (Azouvi *et al.* 2002; Rorden & Karnath 2010; Bonato 2012; Azouvi 2017). In about half of these patients, endogenous plasticity drives neglect recovery during the sub-acute stage (<3 months post-stroke) (Ringman *et al.* 2004; Milner & McIntosh 2005). However, in the other half of the patients, spontaneous recovery is more restricted and neglect persists at the chronic stage. This can limit the overall outcome of rehabilitation by adversely affecting mobility, discharge destination, length of hospital stay and independence in self-care skills (Barer 1990; Jehkonen *et al.* 2006; Patel *et al.* 2007; Di Monaco *et al.* 2011). In light of these functional implications, improving the symptomatology of spatial neglect is an important aim in stroke rehabilitation.

There is currently no clinically established effective treatment of neglect (Bowen *et al.* 2007; Bowen *et al.* 2013). Here, we focus on a promising technique called prism adaptation (PA), introduced by Rossetti *et al.* (1998) in which patients gradually adapt to right-shifting prism glasses through repeated visually-guided pointing movements. In both healthy individuals and neglect patients, prism exposure triggers error-corrective processes that lead to the development of a leftward sensorimotor after-effect (AE) (Chapters 4, 5 & 6; Redding & Wallace 1996; Sarri *et al.* 2008; Aimola *et al.* 2012). However, unlike in healthy individuals, the AE generalises beyond the trained sensorimotor domain in neglect patients, leading to improvements over a broad range of tasks (for review, see Jacquin-Courtois *et al.* 2013). The main limitation of this approach, like most neglect rehabilitation techniques, is its relatively short-lasting effect. Following a single session of PA, improvements have been

reported to last for up to a few days (Rossetti *et al.* 1998; Farnè *et al.* 2002; Pisella *et al.* 2002; Jacquin-Courtois *et al.* 2008). Traditionally, authors have tried to extend the duration of these improvements empirically, by using repeated sessions of PA over consecutive days, aiming to consolidate the cognitive after-effect. With this approach, improvements have been reported to last for up to 4 weeks (Frassinetti *et al.* 2002; Shiraishi *et al.* 2008; Serino *et al.* 2009; Vangkilde & Habekost 2010).

Throughout the work reported in this thesis, we aimed to improve PA-based neglect rehabilitation further, based on a principled mechanistic understanding of sensorimotor memory formation. In Chapters 5 and 6, we investigated the neuro-computational mechanisms underlying AE persistence and its modulation by primary motor cortex (M1) anodal transcranial direct current stimulation (tDCS). We found that, in young and older adults, the level of excitation:inhibition in M1 determines memory persistence so that lower inhibition is associated with a more persistent AE. In both populations, the application of anodal tDCS to M1 during PA enhanced the AE memory trace. In light of these findings, the goal of the present experimental chapter is to test the prediction that repeated sessions of PA combined with left M1 anodal tDCS would enhance retention of the adaptation, stabilising prism AEs from one day to the next in both the sensorimotor (leftward shift in pointing) and cognitive domains (neglect improvement). To do so, we designed a between-subject double-blinded sham-controlled clinical trial and aimed to test 15 neglect patients in each intervention arm (PA + M1 anodal tDCS; PA + M1 sham tDCS). Despite our joint efforts with a clinical team, challenges with patient recruitment demanded us to re-visit the study design to make it a within-subject double-blinded sham-controlled crossover study. Between 2014 and 2017, 182 potential participants were identified but only 6 neglect patients were included in the study (Figure 8.1). Furthermore, some of these patients had a relatively mild neglect and had developed compensatory

strategies that made their assessment with traditional clinical tests rather challenging. The relatively small sample size together with the high level of performance of some patients, limited our ability to address the question of interest (i.e. evaluate the clinical effect of the intervention). Thus, after briefly reporting the clinical results, the present chapter will focus mainly on characterising the effect of M1 anodal tDCS on repeated sessions of PA in neglect patients, and extracting the time course of putative internal computational states engaged during the intervention.

8.2 Material and Methods

8.2.1 Participants and recruitment process




Figure 8.1: Recruitment flow chart.

The inclusion criteria for this study were:

- male or female, aged between 18 and 90 years old;

- at least 3 months post ischaemic or haemorrhagic stroke or brain damage to the right hemisphere;
- diagnosis of left spatial neglect confirmed by evidence of neglect on at least one of the tests in the battery of neglect tests administered during the recruitment interview (conventional Behavioural Inattention test, Wilson *et al.* 1987; apple cancellation, Bickerton *et al.* 2011; computerised cancellation task Dalmaijer *et al.* 2015)

The absolute exclusion criteria were:

- understanding of English insufficient to give informed consent;
- limited verbal communication in the form of dysphasia;
- history of drug abuse;
- history of dementia or any other psychiatric illness;
- history of epilepsy, seizures, febrile convulsions as a child or recurrent fainting fits

Figure 8.1 shows the recruitment flowchart for this study. One hundred eighty-two potentially suitable patients were identified by various means (96 at the Oxford Centre for Enablement, 58 through the Clinical Research Network, 28 via general public advertisement; Figure 8.1). Forty-four patients were screened over the phone to assess their suitability to take part in the study. Spatial neglect was suspected in 25 of them but 4 presented contraindications to brain stimulation, preventing them from taking part in the study. The remaining 21 patients were invited for a recruitment interview. Ten patients showed measurable neglect during this session but 2 of them had severe contraindications to receiving tDCS and 1 patient dropped out. Thus, 7 patients were invited to participate in the study but 1 of them did not show any

measurable neglect on the battery of tests used during the first testing session and was subsequently excluded. In total, only 6 patients participated in the study. Their characteristics are detailed in [Table 8.1](#) and their brain imaging data are presented in [chapter 7](#). Five of them (patients 1-5) received both interventions (PA + M1 anodal tDCS and PA + M1 sham tDCS) in counterbalanced order and one patient (patient 6) received one phase only (PA + M1 sham tDCS).

Patient	Sex	Age	Etiology	C	TPO		Left HA	ARAT		
					<i>anodal</i>	<i>sham</i>		<i>L</i>	<i>R</i>	
1	NT	F	51	right MCA infarct	+	20	32	+	0	57
2	MR	M	77	right ICA infarct	-	96	81	+	2	57
3	SR	M	53	right ICA dissection	-	15	25	+	57	57
4	RW	M	57	right MCA infarct	+	27	33	+	0	57
5	DS	M	62	right MCA infarct	-	30	22	+	0	57
6	UK	M	62	right BG haemorrhage	-	x	4	+	x	x

Table 8.1: Patient characteristics. ARAT = Action Research Arm Test (Lyle 1981); BG = basal ganglia; C = craniectomy; HA = hemianopia; ICA = internal carotid artery; MCA = middle cerebral artery; TPO = time post-onset of stroke (months); L = Left; R = Right; + = yes; - = no; x = not tested.

8.2.2 Experimental design

Neglect patients were tested in a within-subject double-blinded sham-controlled cross-over study including two phases (anodal tDCS phase; sham tDCS phase) that lasted about 12 weeks each. A minimum washout period of 1 month was allowed between the two phases. After the recruitment screening interview, each phase included two pre-intervention testing sessions (pre-test 1 & pre-test 2) in which baseline neuropsychological assessment of neglect (weeks 1 and 3) were taken ([Figure 8.2](#)). Neuroimaging data ([chapter 7](#)) were collected within this pre-intervention period (usually week 2). In week 4, patients underwent 5 consecutive daily sessions of PA combined either with left M1 anodal tDCS or left M1 sham tDCS. Post-intervention follow-up was subsequently conducted over 8 weeks with neuropsychological testing at weeks 5, 7

and 11 (post-tests 1, 2 & 3). A final interview was conducted at the end of the 12 weeks.




Figure 8.2: Experimental design. All patients underwent both experimental phases (PA + M1 anodal tDCS; PA + M1 sham tDCS) with a washout period of at least a month in-between. Two patients received the sham tDCS phase first (P2 & P5) and three patients received the anodal tDCS phase first (P1, P3 & P4). Patient 6 underwent the sham tDCS phase only. Each phase started with an initial interview during which suitability to take part in the study was assessed. Baseline performance on the neglect test battery (NTB) was assessed at weeks 1 and 3. Then, the intervention week (week 4) included 5 consecutive prism adaptation sessions combined with either left M1 anodal tDCS or left M1 sham tDCS. Post-intervention performance on the NTB was assessed at weeks 5, 7 and 11. A final interview was conducted at the end of the 12 weeks.

8.2.3 Assessment of spatial neglect

To assess the severity of spatial neglect, a battery of 7 visuospatial tasks was administered to each patient in every testing session (Figure 8.2). The neglect test battery (NTB) included 3 pen-and-paper tests, 2 computerised tests and 1 naturalistic task. During testing, patients sat at a table in front of the experimenter in a quiet room. Prior to administering each test, the experimenter explained the instructions

and ensured that patients understood them. Tests were always centred on patients' body midline. Patients were asked to perform the tests with their ipsilesional right hand. The experimenter was blinded to the tDCS condition but not to the time of the session respective to the intervention (i.e. pre-intervention or post-intervention).

8.2.3.1 Neglect test battery

The neglect test battery (NTB) included: 1) bells cancellation (Gauthier *et al.* 1989), 2) Ota-Girardi cancellation (also used in O'Shea *et al.* 2017), 3) 3 computerised cancellation tasks (Dalmajer *et al.* 2015), 4) 5 line bisections, and 5) gardening task. All pen-and-paper tests (tests 1, 2, 4) were administered on A3 paper in order to maximise the space spanned by the test.

- **Bells cancellation test.** Patients were instructed to cross out all the bells (in total: 35 bells distributed evenly across 7 columns) as quickly and accurately as possible, and to put their pen down once finished (Gauthier *et al.* 1989) (example performance in [Figure 8.3A](#)). A maximum completion time of 3 minutes was allowed.
- **Ota-Girardi cancellation test.** An A3 stimulus sheet containing 120 squares with a tail (40 in each quadrant) and 36 squares without a tail (9 in each quadrant) was presented to the patients. The instruction was to cross out the tailed squares and circle the squares without a tail (example performance in [Figure 8.3B](#)). Both tasks had to be performed in parallel rather than one after the other. A maximum completion time of 5 minutes was allowed.




Figure 8.3: Example pen-and-paper tests. **A.** Bells cancellation test performed by patient 2 (phase 1, pre-test 1), showing missed items on the left side of the page. **B.** Ota-Girardi cancellation test performed by patient 4 (phase 1, pre-test 1) showing a typical left-to-right trajectory with a lot of omissions on the left part of the page. **C.** Example computerised test, performed by patient 4 (phase 1, pre-test 2). The cancellation heat map superimposed shows clear omission of targets (rooks) in the left side of space. **D.** Line bisection performed by patient 5 (phase 1, pre-test 1) showing a rightward bias. **E.** Gardening task performed by patient 5 (phase 2, pre-test 1), showing omission of the left space.

- **Cancellation tasks.** Computerised cancellation tasks were administered on a 21.5-inch touch screen (Inspiron All-in-One 2320; DELL) positioned about 60 cm away from patients, using the software CancellationTools (Dalmaijer *et al.* 2015). This display allowed us to present a fixed visual array and record the position and timing of patients' touch on the screen (Figure 8.3C). The task was to touch once only all target items as quickly as possible. The target was presented to patients prior to beginning the task. The maximum time allowed was

5 minutes. Every time patients touched the screen either a cross appeared on the screen ('visible' condition), indicating that the patient had already touched this item, or no cross appeared ('invisible' condition), requiring patients to remember the regions of space already covered. By increasing the spatial working memory load, the later condition is thought to provide a more sensitive measure of spatial neglect (Wojciulik *et al.* 2004). In each experimental session, patients performed 2 'invisible' versions of the task and 1 'visible' version, with different visual stimuli for each. The visual arrays consisted in 80 targets evenly distributed across the screen and 440 distractors (10 types of distractor, repeated 44 times each). Parallel versions of the tests were used to minimise general learning effects.

- **Line bisection.** Patients were instructed to put a mark in the middle of a single line presented on a landscape A3 paper aligned to their body midline (example performance in [Figure 8.3D](#)). They were asked to place the mark as quickly as possible, and were not allowed to use their hand or a tool to measure the line. Each of the five lines was presented at a different time point throughout the testing session. No maximum completion time was imposed on this test.
- **Gardening task.** A 26.5×13.5 cm² plastic turf was placed in front of patients and centred on their body midline. Patients were instructed to distribute 12 flowers (one after the other) evenly all over the turf (example performance in [Figure 8.3E](#)). No time limit was imposed on this task.

8.2.3.2 Posner cueing task

The Posner cueing task was as described in Rengachary *et al.* (2009) ([Figure 8.4](#)). Stimuli were displayed on a 17-inch monitor centred on patients' body midline. The display (black background) contained a central fixation cross and two eccentric square




Figure 8.4: Posner cueing task. Time line of a single trial of the Posner cueing task. The onset of the trial was signaled to participant by the fixation cross turning green, followed by an arrow pointing to the left or to the right frame. After a delay randomised between 1000 and 2000 ms, the target (an asterisk) appeared for 300 ms in the cued side on 75% of the trials (valid trials) or in the opposite side on 25% of the trials (invalid trials). An inter-trial interval (ITI) of 2300 ms separated subsequent trials.

frames (in white) positioned along the horizontal meridian to the left and right of fixation. All patients reported being able to see the eccentric square frames when fixating on the central cross hair. Patients were instructed to maintain central fixation on the cross hair at all time during the task and the experimenter monitored compliance to this instruction. The onset of a new trial was signalled by a colour change, from white to green, of the fixation cross. 800 ms later, an arrow cue pointing left or right appeared at fixation for 2360 ms. After a delay, jittered between 1000 to 2000 ms, the visual target (an asterisk) appeared for 300 ms inside one of the two frames. On 75% of the trials, the target appeared at the cued location (valid trials), while on 25% of the trials, it appeared at the opposite location (invalid trials). Patients were instructed to press the space bar (keyboard located in front of them, on their body midline) with their right non-paretic hand as quickly as possible when they detected a target. An inter-trial interval of 2360 ms separated subsequent trials. Patients performed 4 to 5 blocks of 24 trials per testing session. Each block contained 18 valid trials (9 on the left, 9 on the right) and 6 invalid trials (3 on the left, 3 on the right).

Patients were allowed to take a break in between blocks if required.

8.2.4 Intervention procedure

The intervention included 5 consecutive daily sessions of PA combined with anodal or sham left M1 tDCS (Figure 8.2). PA was performed on the manual set-up, using the same PA paradigm as in other chapters (Chapters 4, 5 & 6). Patients sat in a chinrest, in front of a central target and two lateral targets. The distance between their eyes and the central target was approximately 57 cm. The two lateral targets were situated 10 cm to the left and right of the central target. For one patient (Patient 4), a special PA set-up was built so that it could be mounted on his wheelchair. This set-up included a similar pointing board but no chinrest. The reason for this was that the patient's position in the wheelchair prevented him from leaning forward in a chinrest.

Patients performed two types of trial: closed-loop pointing (CLP) and open-loop pointing (OLP). On CLP trials, patients made speeded reaching movements to point with their non-paretic right index finger to the left and right targets, in pseudo-random order. Continuous vision of the hand was available, except for the start position and the first $\approx 30\%$ of the reach trajectory, which were occluded by the chinrest. The removal of visual feedback during the initial phase of the reach movement limited strategic adjustments of hand position (Redding & Wallace 1996; Redding & Wallace 2001; O'Shea *et al.* 2014; O'Shea *et al.* 2017). Both accuracy and speed were emphasised on CLP. At the beginning of prism exposure, CLP became right shifted but patients learnt from error feedback and gradually recovered accuracy. On OLP trials, patients made naturally paced pointing movements at the central target. The target was only visible at the trial onset but was then occluded during the movement. As a result of PA, OLP shifted leftwards, demonstrating the presence of an after-effect.

Every day of the intervention, patients first performed baseline CLP and OLP measures (20 CLP trials followed by 15 OLP trials). Then, the tDCS electrodes were placed on the scalp and stimulation started. The stimulation was on during the entire adaptation phase (20 minutes). The adaptation phase consisted of 6 blocks of prism exposure (100 CLP trials in total; 10 trials in blocks E1-2, 20 trials in blocks E3-6; equal number of trials to each of two target locations) alternating with 6 blocks of AE measurement (90 OLP trials in total; 15 trials per block). At the end of the adaptation phase, the electrodes were removed and patients had a 10-minute rest delay without vision. After this delay, retention of prism AE was assessed (45 OLP trials). At the end of each daily intervention session, patients were asked to report any change in attention, fatigue or pain. No impairment in cognitive state was reported by any patient during the intervention phase.

Patients performed 15 OLP trials in each behavioural testing session to measure long-term retention of the AE.

8.2.5 Transcranial direct current stimulation

The tDCS parameters were as previously described (Chapters 5 & 6; O'Shea *et al.* 2017). Stimulation was delivered by a battery driven DC stimulator (Neuroconn GmbH, Ilmenau, Germany) connected to two 7×5 cm sponge electrodes soaked in a 0.9% saline solution. The electrodes were positioned immediately before stimulation onset and removed as soon as the stimulation ended. The anode electrode was centred over C3 (5 cm lateral to Cz) corresponding to the left primary motor cortex according to the international 10-20 System (Herwig *et al.* 2003). The reference electrode (cathode) was placed over the right supra-orbital ridge. During anodal stimulation, the current intensity was set to 1 mA for 20 minutes with a ramp-up

and ramp-down period of 10 seconds prior and after the stimulation, respectively. During sham stimulation, the current also ramped up and down for 10 sec but no stimulation was delivered during the 20 minutes. Instead, small current pulses (110 μ A over 15 ms) occurred every 550 ms to simulate the tingling sensation associated with real anodal stimulation. Both experimenters and participants were blinded to the stimulation condition.

8.2.6 Neglect test scoring

8.2.6.1 Neglect score

In order to minimise multiple comparisons, all tests included in the NTB were assessed using the same method and averaged into a single neglect score. The common feature across these 5 types of test is that they require patients to indicate the part(s) of space they have successfully engaged with (either by crossing-out items, marking the centre of a line, or distribute flowers). Neglect is characterised by a rightward deviation in the spatial distribution of performance (see [Figure 8.3](#)). To capture this key defining feature of neglect, an equivalent of the centre of cancellation, a measure introduced by Binder and colleagues (Binder *et al.* 1992) and popularised by Rorden and Karnath (Rorden & Karnath 2010), was computed for each sub-test. We call this measure Centre of Performance (CoP) ([Figure 8.5](#)).

For cancellation tests, the CoP is the average horizontal position of all cancelled targets, standardised so that a value of -1 and 1 correspond to the leftmost and rightmost targets respectively. When applied to the line bisection test, the CoP is the average position of the mark over the 5 repetitions of the task, standardised between -1 and 1, corresponding to the left and right end of the line respectively. For the gardening task, the average horizontal position over the 12 flowers was calculated and standardised between -1 and 1, corresponding to the left and right edge of the turf




Figure 8.5: Centre of performance. For each type of test, the centre of performance (CoP) was calculated as the centre of mass of patients' performance and normalised between -1 and 1.

respectively. For each testing session, a single neglect score was obtained by averaging the CoP measure over the 7 sub-tests (i.e. Bells cancellation test, Ota-Girardi cancellation test, 3 computerised cancellation tests, line bisection, gardening task). The advantage of this scoring method is that it captures neglect severity as a continuous attentional gradient and does not rely on an arbitrary division of space (e.g. performance on the left versus performance on the right). The greater this score, the more severe neglect is.

For group mean figures and statistical testing, neglect scores were z-transformed and normalised by the average of the two pre-intervention testing sessions of each experimental phase. This was achieved as follows. For each testing session and each of the 7 sub-tests, the CoP measure was z-transformed using the following formula:

$$zCoP_i = \frac{CoP_i(t) - \overline{CoP_i}}{S} \quad (8.1)$$

where $CoP_i(t)$ is the centre of performance on the sub-test i during the session t , and $\overline{CoP_i}$ and S are the average and standard deviation of the CoP for the sub-test i across all testing sessions of both phases. The z-transformed CoP was then averaged across all sub-tests for each testing sessions to obtain a z-transformed neglect score.

This score represents how exceptionally good or bad performance was during any particular testing session compared to the average patient performance (over the entire experimental phase), and taking into account patients' inter-session variability. A value greater or lower than zero indicates more or less right-shifted performance than usual, respectively. For group mean figures, the z-transformed neglect score averaged over the two pre-intervention testing sessions was subtracted from the z-transformed neglect scores of all testing sessions for each phase and each patient separately. Thus, the resulting group average plot ([Figure 8.7A](#)) represents the change in z-transformed score normalised by baseline. This score was chosen as the primary outcome and compared across the two intervention arms.

8.2.6.2 Posner cueing task

Behavioural performance on the Posner cueing task was assessed by hit rates (HR). To quantify spatial bias in performance, i.e. the core defining clinical feature of spatial neglect (Heilman & Van Den Abell 1980), a lateralisation index (LI) was computed as follows:

$$LI = \frac{(Right - Left)}{(Right + Left)} \quad (8.2)$$

where 'Right' and 'Left' refer to the HR in the respective hemifield. Thus, the LI reflects the degree of performance bias in favour of the right versus left hemifield ([Figure 8.8B](#)). A positive LI denotes a rightward bias, i.e. the tendency to report targets on the right more often than targets on the left. Neglect improvement on the Posner task is therefore defined as a reduction (normalisation) of this this metric. To test the prediction that PA would induce a greater and/or longer-lasting improvement on the Posner task when combined to anodal tDCS compared to sham tDCS, a percentage change in LI was calculated as follows:

$$\% \Delta LI(t) = 100 \times \frac{LI(t) - \overline{LI_{baseline}}}{\overline{LI_{baseline}}} \quad (8.3)$$

where $LI(t)$ is the LI obtained in a post-test session and $\overline{LI_{baseline}}$ is the LI averaged over the two pre-intervention sessions. This score was chosen as the secondary outcome of the intervention.

8.2.7 Analysis of performance during the intervention

Previous research investigating the effect of anodal tDCS on repeated sessions of motor skill learning spaced over multiple days differentiated between total, online (within-day) and offline (between-day) effects (Reis *et al.* 2009). We used a similar approach to characterise the effect of anodal tDCS on the intervention.




Figure 8.6: Schematic representation of the online, offline and total effects.

The total effect was calculated as the magnitude of prism AE obtained at the end of each session (Figure 8.6 in black: average deviation on the 10-minute retention OLP block minus average deviation on the 3 pre-intervention OLP measurements). The online effect was defined as the shift in OLP occurring every day as a result of PA (Figure 8.6 in blue: average deviation on the 10-minute retention OLP block minus average deviation on the baseline OLP of the same day). The offline effect was calculated as the shift in OLP occurring between the end of a PA session and the beginning of the next one (Figure 8.6 in purple: average deviation on the baseline OLP of the n -th day minus average deviation on the 10-minute retention OLP block of the $(n-1)$ -th day). In case of a stimulation-induced increase in total effect, this

approach would enable us to investigate whether it is driven by increased adaptation within-sessions and/or decreased forgetting between-sessions.

8.2.8 Computational modelling of the intervention

The weighted two-state linear time-invariant model described in [chapter 4](#) was modified to fit PA data over multiple consecutive sessions. As before, the model was defined as follows on CLP trials:

$$\begin{aligned}
 e(n) &= bias_{CLP} + d - X(n) \\
 X(n) &= x_f(n) + x_s(n) \\
 x_f(n) &= A_f x_f(n-1) + B_f e(n-1) \\
 x_s(n) &= A_s x_s(n-1) + B_s e(n-1) \\
 0 < A_f < A_s < 1; \quad 0 < B_s < B_f < 1
 \end{aligned} \tag{8.4}$$

where the endpoint error on the n -th trial $e(n)$ is modelled as the sum of a bias term $bias_{CLP}$ and a prism effect d from which is removed the sum of the states of the two systems $X(n)$. The learning and retention rates (B and A) determine the speed of within-session learning and decay of each system. On OLP trials, where prisms are removed and no visual feedback is available, the weighted two-state model therefore becomes:

$$\begin{aligned}
 e(n) &= bias_{OLP} - X(n) \\
 X(n) &= w_f x_f(n) + w_s x_s(n) \\
 x_f(n) &= A_f x_f(n-1) \\
 x_s(n) &= A_s x_s(n-1) \\
 0 < A_f < A_s < 1; \quad 0 < B_s < B_f < 1 \\
 0 < w_f < 5; \quad 0 < w_s < 5
 \end{aligned} \tag{8.5}$$

Equations 8.4 and 8.5 differ in four ways. First, a different bias term is used because of baseline differences in pointing accuracy between CLP and OLP. Second, the prism effect is set to zero to reflect removal of the prisms. Third, the learning rates of both systems (B_f , B_s) are set to zero because the absence of visual feedback prevents error-dependent learning on OLP. Fourth, weighting factors (w_f , w_s) are introduced to allow for a differential contribution of the fast and slow systems to motor behaviour during AE measurements (see chapter 4).

In order to account for the 24 hours separating consecutive PA sessions, two extra parameters (one for each system) were added to the model. They define a linear time-invariant relationship between the state of a system at the end of a PA session $x^{d-1}(n_{max})$, and its state at the beginning of the next session $x^d(n_0)$, as follows

$$\begin{aligned} x_f^d(n_0) &= D_f x_f^{d-1}(n_{max}) \\ x_s^d(n_0) &= D_s x_s^{d-1}(n_{max}) \\ -10 < D_f < 10; \quad -10 < D_s < 10 \end{aligned} \tag{8.6}$$

Conceptually, this means that because we could not record patients' motor behaviour during the 24 hours separating two consecutive intervention sessions, we were agnostic about the fate of the two systems during this time. This simple scaling therefore allows systems to do anything between the following options:

1. continue learning in the same direction ($D > 1$)
2. not decay at all ($D = 1$)
3. decay a little ($0 < D < 1$)
4. decay back to zero ($D = 0$)

5. learn in the opposite direction ($D < 0$)

In the latter case, the sign of the state of a system (i.e. its belief about the direction of the visual disturbance) would flip from one day to the next, potentially creating competing memory traces between systems (see [Figure 8.12](#)). Note that because it is time-invariant the inter-day scaling factors D_f and D_s apply similarly to all four 24-hour breaks of the intervention week (day 1-2, day 2-3, day 3-4, day 4-5). The eight free parameters (A_f , A_s , B_f , B_s , w_f , w_s , D_f , D_s) were fitted to the group mean data, independently for each stimulation condition, using a least-mean-squared-error procedure implemented in MATLAB (The MathWorks Inc., version R2014b).

8.2.9 Statistical analysis

Because of the small sample size, non-parametric testing was used. The Mann-Whitney-Wilcoxon test (Mann & Whitney 1947; Fay & Proschan 2010) and Wilcoxon signed rank test (Wilcoxon 1945) were used as non-parametric alternatives to a two-sample and one-sample t -test, respectively. These tests were implemented in Matlab (The MathWorks inc. version R2014b). Patient 6 was excluded from group statistical analyses contrasting the two experimental phases because this patient received one intervention only. Because of limited statistical power due to small sample size, the statistical results reported in this chapter are not corrected for multiple comparisons.

8.3 Results

8.3.1 Effect of the intervention on clinical measures of spatial neglect

As mentioned in the introduction, due to recruitment challenges, the target sample size could not be met. To date, five patients have undergone both experimental

phases (anodal tDCS and sham tDCS). This sample size is too small to address the question of interest (compare the clinical efficacy of the two interventions on the primary and secondary outcome measures) with sufficient statistical power (Minarik *et al.* 2016). The present section reports patients' performance on the clinical tests of neglect throughout the two phases of the study but care should be taken regarding interpretation of statistical results.

8.3.1.1 Primary outcome: neglect score

Figure 8.7 shows the group mean normalised z-transformed neglect score, as well as individual patients' neglect score throughout the study. On the first testing session of each phase, all patients demonstrated measurable neglect as indicated by a positive neglect score (i.e. rightward bias in performance in line bisection, cancellation and gardening tasks; Figure 8.7B-G). There were no differences in pre-intervention neglect score between the two phases in any of the two pre-test sessions (pre-test 1: $Z = -0.13$, $p = 0.89$; pre-test 2: $Z = -0.13$, $p = 0.89$), indicating successful washout between the two phases (i.e. no carry over of learning effects from one phase to the other). During the pre-intervention period, the mere repetition of the NTB between the two pre-test sessions induced an improvement (i.e. reduction) in neglect score in both phases (Figure 8.7A: Pre 1 vs. Pre 2, Wilcoxon signed rank test; sham tDCS phase: $Z = 1.68$, $p = 0.047$, 1-tail; anodal tDCS phase: $Z = 1.62$, $p = 0.053$, 1-tail). However, the magnitude of this practice effect did not differ between the two phases ($Z = -0.13$, $p = 0.89$).




Figure 8.7: Group mean and individual neglect scores. **A.** The group mean normalised z-transformed neglect score (error bar = SEM) is plotted for all testing sessions of the anodal (in red) and sham (in blue) tDCS phases. Non-parametric Wilcoxon signed rank tests showed no significant difference between the two conditions at any time point. **B-G** Individual neglect scores (error bar = SEM) are plotted for patients 1-6. Asterisks indicate neglect scores significantly greater than zero (one-tailed Wilcoxon signed rank test; **: $p < 0.01$, *: $p < 0.05$, (*): $p < 0.1$). Underneath each plot is represented the patient's study time line (in days). Time is expressed in reference to the intervention on top of the line, and in reference to the first testing session on the bottom.

The key hypothesis was that patients would show greater and/or more sustained response to the intervention in the anodal tDCS phase compared to the sham tDCS phase (O'Shea *et al.* 2017). In [Figure 8.7A](#), this would be manifested as more negative normalised z-transformed neglect score in some or all post-test sessions in the

anodal tDCS phase (in red) compared to the sham tDCS phase (in blue). Statistical analyses however revealed that the normalised z-transformed neglect score obtained in the anodal tDCS phase was not lower than the one obtained in the sham tDCS phase, in any of the 3 post-test sessions (post-test 1: $Z = -1.08$, $p = 0.14$, 1-tail; post-test 2: $Z = 1.89$, $p = 0.97$, 1-tail; post-test 3: $Z = 0.81$, $p = 0.79$, 1-tail). Thus, the data provided no support in favour of the hypothesis.

Efficacy of PA alone could not be evaluated in this study because no sham PA intervention was tested. It is therefore impossible to disentangle any PA-induced improvement from mere practice effect. Nonetheless, an improvement relative to the average of both pre-test sessions was observed in post-tests 1 and 2 during the anodal tDCS phase (both $p < 0.05$, 1-tail), and in post-tests 2 and 3 during the sham tDCS phase (both $p < 0.05$, 1-tail).

8.3.1.2 Secondary outcome: Posner cueing task

Figure 8.8A shows the hit rate of individual patients for targets presented in the left and right hemifields. As expected, patients detected targets in the right hemifield more reliably than in the left, one of the clinical features of neglect (Heilman & Van Den Abell 1980). This was quantified at the individual level by computing a LI reflecting the degree of performance bias in favour of the right versus left visual hemifield (Figure 8.8B). There was a trend towards a greater improvement following PA + anodal tDCS compared to PA + sham tDCS in the first post-test session ($Z = -1.62$, $p = 0.053$, 1-tail; Figure 8.8C). However, subsequent post-intervention time points showed no difference between the two conditions (post-test 2: $Z = 0$, $p = 0.60$; post-test 3: $Z = 0.27$, $p = 0.61$; Figure 8.8). The same analysis run on reaction time showed no significant difference between the two conditions (all $p > 0.18$).




Figure 8.8: Behavioural performance on the Posner cueing task. **A.** Graphs show percentage correct reports (HR) of individuals patients (P1-6) for targets presented in the left and right hemifield, across all testing sessions of both phases (red = anodal; blue = sham). Performance in the left hemifield was consistently worse (i.e. lower HR) than in the right. **B.** This was quantified as a lateralisation index (LI), plotted in this panel across all testing sessions of both phases (red = anodal; blue = sham). **C.** Group mean percentage change in LI (error bar = SEM). One-tailed Wilcoxon signed rank tests performed at each time point revealed a trend towards significantly greater improvement (i.e. reduction in the LI) in the anodal tDCS condition compared to the sham tDCS condition at the first post-intervention time point ($Z = -1.62$, $p = 0.053$). Percentage change in LI did not differ between the two conditions in other post-intervention testing sessions (both $p > 0.60$).

8.3.2 Patients adapted to prisms and demonstrated a sensori-motor after-effect




Figure 8.9: Prism adaptation profile. Graph shows patients’ PA data on the first intervention session of the sham tDCS phase. The x-axis represents trial number and the y-axis pointing accuracy relative to baseline (lateral error, in cm, rightward errors are positive). Black wedges indicate blocks throughout which 10° right-shifting prisms are worn. Participants received visual feedback on CLP trials (shaded dark grey), allowing error-dependent learning to occur. By contrast, no visual error feedback was provided on OLP (white background), allowing the measurement of prism AE. Patients demonstrated significant error reduction on CLP and leftward shift on OLP (AE). Neglect patients demonstrated the ability to correct pointing error on CLP (E1 vs. E6: $Z = 1.89$, $p = 0.03$, 1-tail), which subsequently led to a measurable after-effect on OLP (AE1-7: all $p < 0.02$, 1-tail).

Pointing data taken from the first PA session of the sham phase were analysed to determine whether patients could successfully adapt to 10° right-shifting prism glasses (Figure 8.9). The introduction of the visual displacement resulted in rightward errors on the first block of prism exposure (E1: $Z = 2.10$, $p = 0.02$, 1-tail). The magnitude of these errors was small in comparison to healthy individuals, potentially because patients’ movements were slower and contained a lot of in-flight error correction. Patients demonstrated the ability to correct CLP endpoint errors throughout prism exposure (E1 vs. E6: $Z = 1.89$, $p = 0.03$, 1-tail), indicating that they could learn from performance error. As a result, OLP accuracy shifted leftward on the first block

of AE measurement (AE1: $Z = -2.10$, $p = 0.02$, 1-tail) and remained left-shifted on all subsequent AE measurements (AE2-6 and retention block: all $p < 0.02$, 1-tail). No significant retention of the AE was measured the day after ($Z = -0.21$, $p = 0.42$, 1-tail). This pattern of behavioural performance indicates that patients could successfully correct pointing errors during prism exposure, which led to a sensorimotor AE on OLP that was still present 10 minutes after PA but not 24 hours later. Patients' behaviour on this task was therefore qualitatively comparable to those of healthy adults (Chapters 5 & 6). This is an important quality check because the induction of a sensorimotor AE appears to be necessary for neglect improvements to occur after PA (Striemer & Danckert 2010; Rode *et al.* 2015). Patients' data could not be compared directly to those of the age-matched healthy controls tested in [chapter 6](#) because they were collected on a different experimental set-up.

8.3.3 Effect of left M1 anodal tDCS on prism adaptation in patients




Figure 8.10: Time course of pointing accuracy throughout the study. Block mean OLP endpoint error normalised by patients' pre-intervention deviation (group mean, error bar = SEM) is shown for all testing sessions and for the first and last block of each intervention day (red = anodal tDCS condition; blue = sham tDCS condition). Patients shifted leftward in both conditions but there was no difference in pointing accuracy between the two conditions in any of the time points.

Figure 8.10 shows the group mean AE measured during each behavioural testing-session and at the beginning and end of each intervention session. Data are normalised by the average of the 3 pre-intervention OLP measurements (i.e. pre-test 1, pre-test 2, baseline OLP of the first intervention day). This plot shows that patients shifted leftward on OLP as a result of each PA session. The *a priori* hypothesis was that left M1 anodal tDCS + PA would enhance AE persistence, resulting in gradually more left-shifted OLP behaviour throughout the intervention week. To test this hypothesis, one-tailed one-sample Wilcoxon signed rank tests were run separately for days 2 to 5 to evaluate whether the mean error on the baseline OLP block (providing a measure of retention of all preceding PA sessions) were significantly more left shifted in the anodal compare to sham tDCS condition. None of these tests were significant (day 2: $Z = -0.81$, $p = 0.21$; day 3: $Z = -0.54$, $p = 0.29$; day 4: $Z = -0.91$, $p = 0.18$; day 5: $Z = 0.27$, $p = 0.61$; Figure 8.10). The same pattern was found if testing for difference in AE at the end of PA sessions (all $p > 0.10$). The hypothesis was also not confirmed in subsequent post-intervention testing sessions (post-test 1: $Z = 1.09$, $p = 0.86$; post-test 2: $Z = 1.62$, $p = 0.95$; post-test 3: $Z = 0.91$, $p = 0.82$; Figure 8.10). Given the small trend towards greater AEs during the intervention in the anodal tDCS condition, it is unclear whether this non-result simply reflects inadequate sample size or a true null effect of stimulation. Thus, the rest of this section aims to characterise further the effect of anodal tDCS by disentangling potential effects on within-session learning and between-session consolidation/decay throughout the week (Reis *et al.* 2009).

Observation of the temporal evolution of the AE throughout the intervention week revealed inter-session forgetting in the sham tDCS condition, i.e. a decrease in the magnitude of the AE between the end of an intervention day and the beginning of the following one (Figure 8.10: dashed blue line). The profile observed in the anodal

condition however seemed to be characterised by less inter-session forgetting between days 1 to 4. In order to formally assess intra-session adaptation and inter-session forgetting, the total, online and offline effects were computed as described in section 8.2.7 (Figure 8.11).




Figure 8.11: Online (within-day) and offline (between-day) effects of a-tDCS. **A.** Graphs show the total (in black), online (in blue) and offline (in purple) effects for each patient (P1-6) during the anodal tDCS (red background) and sham tDCS (blue background) intervention. **B.** Graphs shows the group mean (error bar = SEM) online effect (leftward shift in OLP accuracy) for each intervention day in the anodal and sham tDCS condition. Anodal tDCS resulted in a smaller online effect on day 2 and 3 (both $p = 0.04$) and a marginally significant greater online effect on day 4 ($p = 0.07$). **C.** Graphs shows the group mean (error bar = SEM) offline effect (rightward shift in OLP) between intervention days in the anodal and sham tDCS condition. Anodal tDCS resulted in significantly lower decay of the AE between days 2 and 3 ($p = 0.04$). Asterisks indicate significant 2-tail one sample Wilcoxon signed rank test (*: $p < 0.05$; (*): $P < 0.10$, not corrected for multiple comparisons).

Two-tailed one sample Wilcoxon signed rank tests were run separately for each day to compare the online and offline effects across the two interventions. Anodal tDCS resulted in significantly smaller online effects (smaller leftward shift in OLP) on

day 2 and 3 (day 2: $Z = 2.02$, $p = 0.04$; day 3: $Z = 2.02$, $p = 0.04$; [Figure 8.11B](#)) but slightly greater online effect (greater leftward shift in OLP) on day 4 ($Z = -1.82$, $p = 0.07$; [Figure 8.11B](#)). The offline effect was reduced between day 2 and 3 in the anodal compared to sham tDCS condition ($Z = -2.02$, $p = 0.04$; [Figure 8.11C](#)). In other words, although anodal tDCS did not affect the overall magnitude of AE at the week level (no change in total effect; [Figure 8.10](#)), it stabilised the AE earlier during the week by reducing inter-session forgetting, subsequently leading to more stable OLP behaviour (i.e. apparently smaller online effect) during early sessions. Potentially due to an fatigue effect accumulating throughout the week or a ceiling effect reached on day 4, the positive effect of anodal tDCS on AE stability was not visible on the last intervention day ([Figure 8.11](#)).

8.3.4 Extracting hidden computational states during the intervention

In order to investigate putative hidden computational states driving behaviour during the intervention week, a multi-day version of the weighted two-state model was fitted to the data. Because of 1) the noisy nature of multi-day PA data in patients, and 2) the small sample size, the model was fitted to the group mean data only and not to individual datasets. Results of the fit are discussed below.




Figure 8.12: Fit of the multi-day weighted two-state model to the intervention PA data. Graphs show the time course of pointing accuracy throughout the 5 intervention days in the anodal (A) and sham (B) tDCS condition. Overlaid are shown the model fit (in red) and state of the two systems (green = fast system; blue = slow system).

The model yielded a good fit to the data for both interventions (PA + anodal tDCS: $R^2 = 0.95$; PA + sham tDCS: $R^2 = 0.96$; Figure 8.12). The inter-day scaling factor D was negative for the fast system (anodal tDCS: $D_f = -5.48$; sham tDCS: $D_f = -4.61$) and positive and close to 1 for the slow system (anodal tDCS: $D_s = 0.91$;

sham tDCS: $D_s = 1.01$) in both conditions. This illustrates little to no forgetting of the slow system from one day to the next, coherent with its high within-session retention factor (anodal tDCS: $A_s = 1.00$; sham tDCS: $A_f = 0.99$). By contrast, the state of the fast system flipped and grew in the other direction from one day to the next. It was as if, during the 24-hour breaks, the fast system learnt to almost completely counteract the slow system. This was particularly visible in the 'baseline' pointing behaviour of later sessions (Figure 8.12: zoom on the baseline CLP and OLP blocks of day 5 in the sham tDCS condition). On the first baseline CLP trial, patients were slightly left shifted, reflecting the fact that the (leftward) adaptation of the slow system was greater than the (rightward) adaptation of the fast system. Within 4 trials however, basal CLP accuracy was restored, due to both systems learning from these performance errors (active de-adaptation). Removal of vision during the baseline OLP block immediately following the CLP block induced spontaneous recovery, explained by a faster (leftward) forgetting of the fast system than the slow system that was still in an leftward-adapted state. This is an interesting phenomenon indicating that, although behaviour appears relatively unchanged on CLP relative to before the intervention, it may be driven by a new equilibrium reached between opposing competing memory traces, a fast-learning fast-forgetting right-shifting one and a slow-learning slow-forgetting left-shifting one. It is tempting to hypothesise that it is this persistent slow system-driven leftward shift that may be responsible for neglect improvements following PA. Formally testing this prediction however would require a larger sample size than we have in this study in order to run correlation analyses.

Finally, the model parameters were compared across the two conditions. The model fit yielded a greater weight of the fast system (anodal tDCS: $w_f = 0.96$; sham tDCS: $w_f = 0.61$) and a lower weight of the slow system (anodal tDCS: $w_s = 0.69$;

sham tDCS: $w_s = 0.80$) in the anodal tDCS condition compared to sham. This resulted in a reduction in relative weight (anodal tDCS: $R_w = 0.42$; sham tDCS: $R_w = 0.57$). This is consistent with the results in age-matched healthy adults ([chapter 6](#)) where R_w was found to be positively correlated with M1 GABAergic inhibition and was found to be reduced by left M1 anodal tDCS during PA.

8.4 Discussion

This study aimed to test the hypothesis that repeated sessions of PA combined with left M1 anodal tDCS would lead to a greater and/or more persistent improvement of neglect symptoms than the same intervention combined with sham tDCS. This prediction derives from the work described in earlier chapters which demonstrated a beneficial effect of M1 anodal tDCS on AE persistence in young and older adults (Chapters 5 & 6; O'Shea *et al.* 2017). Difficulties with patient recruitment forced us to revisit the original study design and to include patients with mild neglect (patients 3 and 6). This, together with the small sample size, limited our ability to address the question of interest, i.e. compare the clinical efficacy of the two interventions. During the intervention week, the data were in favour of an earlier stabilisation of the AE (i.e. reduced inter-day forgetting; [Figure 8.11](#)) in the anodal tDCS condition compared to sham tDCS. However, this did not result in greater persistence of the AE at the post-intervention phase ([Figure 8.10](#)). Consequently, the behavioural data (NTB and Posner cueing task) revealed improvements in neglect measures after the intervention but no clear stimulation-induced benefit ([Figure 8.7 & 8.8](#)).

Previous work from the laboratory showed that excitatory a-tDCS of the left sensorimotor cortex during a single session of PA could induce a long-lasting therapeutic response in 3 chronic spatial neglect patients that were otherwise treatment resistant

(O'Shea *et al.* 2017). However, in the present study, no such benefit was observed. What factors may be responsible for this discrepancy? The main design difference between the proof-of-concept study and the study presented in this chapter is the use of a single session of PA versus repeated sessions here. It might be the case that the effect of an adjuvant excitatory stimulation favourising long-term retention of the AE is more likely to be detectable with a light sub-optimal PA training (1 session) compared to an intensive 1-week training during which AE may plateau regardless of the stimulation regime.

All patients demonstrated leftward sensorimotor after-effects (on OLP) as a result of PA. Consistent with previous studies (Frassinetti *et al.* 2002; Turton *et al.* 2010; Rode *et al.* 2015), leftward shift in pointing behaviour was particularly evident during the first few days of the intervention and the AE subsequently stabilised over the following days. Results of the modelling work suggest that this consolidation effect may be sub-served by competing internal models: a versatile rightward fast system counteracting a stable leftward slow system. This observation explains the presence of 1) spontaneous recovery (i.e. gradual re-emergence of a leftward bias on OLP in the 'baseline' block of subsequent PA sessions) and 2) savings (i.e. faster re-adaptation during the first few trials of prism exposure of subsequent PA sessions) (Krakauer & Shadmehr 2006; Smith *et al.* 2006a; Joiner & Smith 2008; Shadmehr *et al.* 2010). In this scenario, the fast system acts as a suppressor of the slow system, which enables accurate pointing when vision is provided (CLP) but persistence of an after-effect when it is not.

No previous experimental or modelling work has considered the role of the fast and slow systems in the neglect improvements reported following PA. The small sample size of the present study prevented us from tackling this question, but we set out

a method that could be applied to bigger clinical studies. We speculate that the persistence of a leftward memory trace within the slow system might be involved in the beneficial effects of PA for neglect.

8.5 Conclusions

The methodological challenges encountered in this study (i.e. small sample size, high-performing patients) limited our ability to draw definite conclusions. All patients tested showed a good tolerance for the intervention, confirming the feasibility of the use of anodal tDCS as an adjuvant to prism adaptation training in this clinical population. However, no clear clinical benefit was observed in any of the patients in the present study.

Chapter 9

Conclusions and future directions

This chapter outlines the main findings reported in this thesis, and assesses the overall conclusions that can be drawn from the data therein. The limitations, additional questions and future work that might be warranted to address such limitations and questions are also discussed.

9.1 Aims of the thesis

The primary aim of this thesis was to investigate the neuro-computational mechanisms underlying human sensorimotor memory formation during prism adaptation (PA) and its modulation by excitatory anodal transcranial direct current stimulation (a-tDCS) in young and older adults (Chapters 4, 5 & 6). This fundamental work was largely motivated by its clinical relevance to stroke survivors suffering from spatial neglect (Heilman & Van Den Abell 1980; Mesulam 1981; Heilman *et al.* 2002). In this clinical population, PA has been shown to induce therapeutic improvements that generalise beyond the restricted context of the training (Striemer & Danckert 2010; Barrett *et al.* 2012; Newport & Schenk 2012; Jacquin-Courtois *et al.* 2013) but typically decay within a few days (Rossetti *et al.* 1998; Farnè *et al.* 2002; Frassinetti *et al.* 2002). By better understanding the neural mechanisms underlying sensorimo-

tor memory formation during PA in the healthy brain and how to influence them in favour of long-term retention, we were hoping to improve the efficacy of neglect rehabilitation by PA. For this purpose, a variety of neuroimaging (MR spectroscopy, diffusion-weighted imaging, task-based and resting functional MR imaging), genetic and brain stimulation techniques (a-tDCS) were employed in young adults, older adults and stroke survivors, together with behavioural and computational markers of PA. The studies reported in this thesis provide an example of 'bench-to-bedside' translational research, grounded within a principled quantitative theoretical framework (state-space modelling).

9.2 Summary of key findings

The traditional theoretical framework used in the literature on PA posits 2 putative adaptive processes (strategic control and spatial realignment) thought to contribute differentially to motor behaviour during error correction (prism exposure) and AE measurements (after removal of the prisms) (Rossetti *et al.* 1993; Redding & Wallace 1996; Redding & Wallace 2001; Redding *et al.* 2005; Redding & Wallace 2006b; Aimola *et al.* 2012; O'Shea *et al.* 2014). Although it accounts for several behavioural dissociations observed in healthy individuals and neurological patients, the main limitation of this framework is that it lacks precision, and does not allow for quantitative estimation of the processes involved (chapter 1; Petitet *et al.* 2017). To remedy these limitations, we used an alternative theoretical framework (state-space modelling) that also posits the experience of motor performance errors (i.e. rightward pointing error while wearing right-shifting prisms) to recruit multiple processes learning in parallel (Smith *et al.* 2006a). By definition, these processes are hypothesised to exhibit different error-sensitivity and retention properties, and are typically approximated by two systems: one fast that is very responsive to performance error but displays weak

retention, and one slow that learns more slowly but retains over a longer period of time (Smith *et al.* 2006a; Lee & Schweighofer 2009; Inoue *et al.* 2015; Albert & Shadmehr 2017). This theoretical framework relies on the assumption that the learning and forgetting time courses of individual systems are not directly observable in the raw behavioural data (endpoint errors) but can be inferred (quantitatively) from their collective influence on motor behaviour. Conceptually, this framework is compatible with the traditional dual-process framework¹ and has the advantages of expressing in clear mathematical terms the learning and forgetting algorithms involved, thus providing a way of quantifying the contribution of putative processes to behaviour. The PA paradigm used in the studies presented in this thesis was designed to enable a reliable extraction of the hidden computational states contributing to PA and AEs by providing a measure of the evolving dynamics of prism AEs at various stages throughout prism exposure (Inoue *et al.* 2015; O'Shea *et al.* 2017).

In [chapter 4](#), we extended the original two-state model of sensorimotor adaptation (Smith *et al.* 2006a) to allow for a variable contribution of each system to motor behaviour during AE measurements. The rationale for doing so was two-fold. First, there are several reports in the literature of data in which an experimental intervention or pathological process has no impact on motor behaviour during error-dependent learning (i.e. exposure to the perturbation) but affects performance on subsequent AE measurements (i.e. after removal of the perturbation) (Fernandez-Ruiz *et al.* 2003; Hadipour-Niktarash *et al.* 2007; Galea *et al.* 2015; Nemanich & Earhart 2015; Leow *et al.* 2016; O'Shea *et al.* 2017), or vice versa (Buch *et al.* 2003; Galea *et al.* 2010; Aimola *et al.* 2012; Panouillères *et al.* 2015). These observations illustrate a certain de-coupling of behavioural performance between the two experimental con-

¹In fact, a recent study suggested a direct mapping between fast and slow systems and 'explicit' and 'implicit' processes (an equivalent of the 'strategic control' and 'spatial realignment' components of PA) during a visuomotor learning task (McDougle *et al.* 2015).

ditions (error-dependent learning vs. AE measurement), a degree of freedom that cannot be explained with a non-weighted two-state model (Smith *et al.* 2006a). Second, from a theoretical perspective, it is evident that the internal model formed during adaptation to a perturbation (Shadmehr *et al.* 2010; Wolpert *et al.* 2011) should not only incorporate an estimate of the magnitude of the perturbation (e.g. magnitude of prism-induced visual displacement) but should also represent its most likely time course (Körding *et al.* 2007) and context specificity (Vetter & Wolpert 2000; Kluzik *et al.* 2008; Carroll *et al.* 2014). We proposed that the latter two are reflected in the extent to which the state of different learning processes (i.e. fast and slow systems) transfers across the error-dependent learning and AE measurement phases. In practice, we added two extra free parameters to the original model that define a linear scaling of the contribution of fast and slow systems to trials in which the prismatic shift is removed and visual feedback deprived. Thus, the AE can now be conceptualised as a weighted sum of memory traces originating from learning processes that forget at different rates, one short-lasting (fast system) and the other longer-lasting (slow system). This computational concept can be quantified as a relative weight (R_w). We showed in Chapters 4 & 5 that greater relative contribution of the slow system to open-loop pointing trials during PA predicts greater and more stable AEs at retention in young healthy adults.

In chapter 5, we tested and confirmed the prediction that greater M1 excitation:inhibition (E:I) in primary motor cortex (M1) during PA causally increases the relative contribution of the slow system to prism AEs, which in return, is predictive of a more stable memory trace. To test this, we applied M1 anodal or sham tDCS either before or during PA in 24 young healthy male individuals in a within-subject design. Individual trait markers of neurochemical response to M1 a-tDCS (Stagg *et al.* 2009; Kim *et al.* 2014; Bachtiar *et al.* 2015; Barron *et al.* 2016; Antonenko

et al. 2017), white matter microstructure and resting functional connectivity were also acquired in 2 separate ultra-high field MRI scanning sessions. Neuroimaging identified structural, functional and neurochemical covariates of the relative contribution of the slow system to prism AEs. To establish causality, the effect of excitatory a-tDCS was assessed on all three levels (neurochemical, computational, behavioural). This resulted in correlated increases in excitation:inhibition, relative contribution of the slow system, and retention. These relations did not hold, and memory was unchanged, with stimulation applied before adaptation. Finally, we showed that the cognitive state dependent effect was modulated by the BDNF val66met genetic polymorphism. Memory was enhanced by stimulation in Val/Val carriers only, implicating a mechanistic role of activity-dependent BDNF secretion and synaptic long-term potentiation (LTP). In summary, the main conclusion of this study was that sensorimotor cortex E:I (and in particular GABAergic inhibition) causally determines the time constant of memory persistence, explaining phenotypic variation in adaptation decay in young healthy adults.

The findings reported in [chapter 5](#) have important clinical implications for neglect rehabilitation because we identified a neuro-computational mechanism (M1 E:I ratio and relative weight) underlying sensorimotor memory formation, together with a tool (a-tDCS) to modulate it in favour of long-term retention (O'Shea *et al.* 2017). To address potential translatability to post-stroke neglect patients, a population that is typically older than the experimental sample tested in [chapter 5](#), we asked in the next chapter ([chapter 6](#)) whether this mechanism was preserved in the ageing brain. We found that, across individuals, older age was associated with lower M1 GABA and greater long-term retention of the prism AE. A mediation analysis provided cross-sectional evidence that the age-related decline in M1 GABAergic inhibition explained the increased memory persistence, a phenomenon that is somewhat comparable to

the one induced artificially in young adults with a-tDCS. In older participants however, the computational mechanisms by which M1 GABA relates to AE persistence appeared to be different. There was an inverse relationship between M1 E:I ratio and R_w , but the latter did not predict properties of the prism AE at retention. Next, preliminary results of a brain stimulation experiment suggested that it might be possible to increase persistence of the AE memory trace further in older adults by applying a-tDCS to M1 during PA. Baseline M1 E:I ratio was found to constrain this behavioural effect, likely reflecting limited neurochemical effect of a-tDCS in individuals with low baseline M1 GABAergic inhibition.

Data reported in the last two experimental chapters (Chapters 7 & 8) were collected in 6 stroke survivors who participated in a within-subject double-blinded sham-controlled crossover study aiming to compare the clinical efficacy of repeated sessions of PA combined with real versus sham left M1 a-tDCS. Chapter 7 presents the neuroimaging data of these individuals and compares them to those of a group of 23 age-matched healthy controls. Chapter 8 presents patients' performance on clinical tests of neglect severity throughout the study. The neuroimaging results indicate altered inter-hemispheric competition within the lateral occipital cortex, a key region of the dorsal attention network (Fox *et al.* 2006; Shulman *et al.* 2009; Shulman *et al.* 2010), during deployment of bilateral attention in neglect. Among the functional differences exhibited by patients was a pathological hypo-activation within the left posterior superior cerebellum, a region that was distant from the lesion in all patients. We hypothesised that the physiological mechanisms underlying therapeutic improvements following PA in neglect patients might implicate a restoration of normal activation levels within this non-motor cerebellar region. The behavioural results reported in chapter 8 failed to support the prediction that PA combined with left M1 a-tDCS would induce greater and/or longer-lasting improvement of neglect

symptomatology. This null-result might reflect limited statistical power due to small sample size or may indicate a true absence of tDCS-induced benefit in the context of repeated PA sessions. This will be discussed in the following section.

9.3 Limitations and future directions

9.3.1 Does M1 a-tDCS enhance clinical efficacy of prism adaptation for neglect rehabilitation?

The clinical study reported in this thesis ([chapter 8](#)) suffered from two main limitations. First, the sample size was lower than expected, which forced us to reconsider the study design (within-individual design instead of the originally planned between-individual design) and limited our ability to address statistically the question of interest (i.e. compare the clinical efficacy of the two interventions). Second, patients with relatively high baseline performance level were included. It was clear that some patients had developed compensatory strategies that enabled them to perform 'artificially' well on the neglect test battery, at the cost of a high cognitive effort (Azouvi *et al.* 2002). The contribution of these compensatory strategies to behavioural performance is difficult to evaluate and may reduce sensitivity to detect a clinical improvement. For example, it might be the case that the overall performance is unchanged following the intervention in these patients but that it relies less on conscious, effortful strategies, which may translate to meaningful clinical improvements in an ecological context. Thus, we suggest that future research should consider limiting and/or objectively assessing the contribution of compensatory strategies to their primary outcome measure.

Current evaluation of functional impairment relies mainly on the Catherine Bergego Scale (CBS), a list of 10 questions aiming to assess performance on activities of daily

living (Bergego *et al.* 1995; Azouvi 1996; Azouvi *et al.* 2003). Although it provides a sensitive measure of neglect and anosognosia (Chen *et al.* 2012), the CBS may not be appropriate to monitor neglect improvement longitudinally and it relies entirely on subjective evaluation. Instead, the use of ecological tests that reduce patients' awareness of the experimental setting (Azouvi 2017) might provide more valid and clinically relevant measures. Finally, because they recruit additional cognitive processing resources, dual-tasks might offer a way to limit compensatory strategies and reveal deficits in apparently high-performing patients (Husain *et al.* 1997; Bonato *et al.* 2013).

9.3.2 Towards a model of prism adaptation that spans Marr's three levels of analysis

The levels of understanding of a complex system proposed by Marr (1982) provide a useful framework to anyone trying to understand the link between brain and behaviour (i.e. most neuroscientists), and help evaluate the current scientific knowledge, its limitations, and the remaining unanswered questions. The central point made by Marr (who originally formulated his thesis as a critique of reductionist work in neurophysiology) is that understanding a complex system is not the same as just describing it or knowing how to intervene to change it (Krakauer *et al.* 2017). He proposed to subdivide the research effort into three levels of analysis, as follows:

1. Computation (goal)
2. Algorithmic realisation (software)
3. Physiological implementation (hardware)

The *computational* level corresponds to the goal of the behaviour. The *algorithmic* level should identify the set of (mathematical) rules and operations (functions) required to meet this goal. The *implementation* level should give an account of the

neurophysiological mechanisms responsible for the execution of the identified algorithms. Each level should be seen as the realisation of the level before it. Thus, the implementation level (level 3) specifies physiological mechanisms that carry out algorithms (level 2) responsible for solving general computational problems (level 1). It does not mean, however, that investigation should proceed from level 3 to 1. In fact, Marr and others have claimed the contrary: it is easier to infer the algorithm from a principled computational analysis of the problem the nervous system is trying to solve (level 1) than it is from the neural hardware itself (level 3) (Marr 1982; Krakauer *et al.* 2017).

The modelling work carried out in this thesis (weighted two-state model) was at the interface between Marr’s level 1 and 2. Nevertheless, we have tried to relate some aspects of the putative algorithmic level (i.e. the relative contribution of the slow versus fast system to prism AEs, R_w) to biomarkers of its potential neurophysiological implementation (level 3). Our work points towards a critical role of LTP-like synaptic plasticity in primary motor cortex, known to be modulated by GABA (Trepel & Racine 2000) and BDNF (Fritsch *et al.* 2010), in determining the relative contribution of different learning processes to motor behaviour after adaptation.

Recent fMRI work suggests that the time courses of hidden internal computational states inferred with state-space modelling are compatible with the pattern of neural activation observed during sensorimotor adaptation (Kim *et al.* 2015). This result provides another potential link between Marr’s level 2 and 3; that is, the state of different learning processes engaged during adaptation might be reflected in the level of activation of certain cortical and subcortical areas, compatible with the idea that different brain networks process information at different timescales (Fusi *et al.* 2005; Bernacchia *et al.* 2011; Chaudhuri *et al.* 2015). Thus, while the state of the fast and

slow systems might be computed by multiple, distributed neural networks, the data presented in this thesis suggest that their contribution to motor behaviour might depend causally upon the synaptic weight of these networks with M1 (which is at least partly determined by local GABAergic inhibition and activity-dependent release of BDNF). It might, therefore, become possible to build a model that covers all explanatory levels, from neuronal implementation to behavioural expression, through algorithmic realisation. Such a model would need to incorporate multiple neural networks learning (at different rates) an internal model of the perturbation from visual error feedback, and feeding it to M1 with a certain synaptic weight determined by the level of local inhibition.

9.3.3 A better characterisation of the fast and slow systems is needed

Despite the several advantages of the modelling framework used throughout this thesis over the traditional theoretical framework (see [chapter 1](#)), state-space models present one main limitation, that is they are agnostic about the information content of the extracted computational states. In other words, although the two-rate dynamics of sensorimotor adaptation behaviour is well established (Smith *et al.* 2006a; Ethier *et al.* 2008b; Joiner & Smith 2008; Lee & Schweighofer 2009; Vaswani & Shadmehr 2013; Inoue *et al.* 2015; Albert & Shadmehr 2017; Broucke 2017), the psychological underpinning of the fast and slow systems remains elusive (Yavari *et al.* 2013): do they reflect the explicit versus implicit components of adaptation (Taylor *et al.* 2014; McDougle *et al.* 2015)? Proprioceptive versus visual internal model updating ([Appendix A](#))? Model-based versus model-free learning (Wolpert *et al.* 2011)? Environmental versus body-based learning (Berniker & Kording 2008)? In-flight versus inter-trial corrections (O’Shea *et al.* 2014)? Something else (Broucke 2017)? A better characterisation of the fast and slow systems is likely to contribute elucidating some

unanswered questions of this thesis such as 1) why is the relationship between M1 E:I and R_w inverted in older adults compared to young individuals?, or 2) what is the computational basis of neglect improvements following PA?

Future research should investigate the aforementioned questions. We propose, for example, that interleaved measures of modality-specific AEs (e.g. visual straight-ahead judgement, proprioceptive straight-ahead pointing; see [chapter 1](#)) in lieu of the OLP blocks used in the experiments reported in this thesis (see PA protocol; [chapter 4: Figure 4.1](#)) would provide some insight on the dynamics of visual and proprioceptive recalibration during PA (Harris 1963; Redding & Wallace 1996; Redding & Wallace 2001; Hatada *et al.* 2006c). One may hypothesise fast proprioceptive recalibration due to low confidence in proprioceptive estimates, and slower but more durable visual recalibration as prediction errors get correctly assigned to the visual domain (Petitet *et al.* 2017). Compared to simply considering the endpoint error, kinematic recording of reaching movements (O’Shea *et al.* 2014) might also provide richer information on the signals driving learning in these two systems.

Some participants in the studies reported in this thesis exhibited spontaneous recovery in the absence of active washout, i.e. they shifted in the direction opposite to the visual displacement (i.e. leftward) within blocks of AE measurements. This pattern of performance cannot be accounted for by two-state models because behavioural changes within blocks of AE measurements are only explained in terms of either 1) active learning of the absence of perturbation, 2) active disengagement of the memory trace, or 3) passive regression towards the non-adapted state (Vaswani & Shadmehr 2013). In practice, this means that, with the current architecture of state-space models, at most motor behaviour remains stable on AE measures (i.e. slope = 0), but it cannot evolve towards a more adapted state. The fact that it does in certain

participants might reflect the ability of the sensorimotor system to *extrapolate*, i.e. continue learning even after removal of the teaching signal (visual feedback) (Guigon 2004). To our knowledge this is an understudied phenomenon that challenges current state-space models and therefore merits more attention.

9.3.4 Explaining inter-individual differences in sensitivity to brain stimulation

Our magnetic resonance spectroscopy (MRS) and genotyping results (Chapters 5 & 6) suggest that physiological markers can be used not only to interrogate the mechanisms by which tDCS alters behaviour, but also to assess *a priori* the likelihood of an individual to respond to brain stimulation. The latter point is crucial to individualise brain stimulation interventions and optimise their efficacy (Rossi *et al.* 2009; Sandrini *et al.* 2011; Bestmann & Walsh 2017). For example, based on our results, pre-treatment of individuals with α -linolenic acids (Hadjighassem *et al.* 2015) or pre-biotics (Savignac *et al.* 2013), two factors known to elevate BDNF levels, might promote the online M1 a-tDCS effect reported in chapter 5, especially in Met carriers who show reduced activity-dependent BDNF secretion (Egan *et al.* 2003; Chen *et al.* 2005).

Although it appears to partly determine behavioural response to M1 a-tDCS (chapter 5 & 6; O'Shea *et al.* 2017), inter-individual variability in the neurochemical response to stimulation (as measured by MRS) remains largely unexplained. Little is known about the factors influencing the effect of tDCS on GABA and glutamate concentration. Current evidence suggests that opposite stimulation polarities (anodal versus cathodal M1 tDCS) have opposite effects on sensorimotor GABA concentration (Stagg *et al.* 2009; Antonenko *et al.* 2017). Because the spatial distribution of the stimulation-induced intra-cranial electric field is non uniform and may depend on the

gyral orientation relative to the main current direction (Datta *et al.* 2009; Rahman *et al.* 2013), it is not unreasonable to speculate that the same tDCS montage may have opposite effects on sensorimotor GABA depending on the underlying anatomy of different individuals. It is unclear however, which brain regions are the key causal nodes in that respect – given the spatial diffusivity of the electric field induced by conventional M1 tDCS (Miranda *et al.* 2006), the reduction of GABA in M1 might even involve excitation of a distant inter-connected region. We would argue that the development of subject-specific models of stimulation-induced electric field (e.g. SIMNIBS; <http://simnibs.de/version2/start>) and their validation with *in vivo* physiological recording (Huang *et al.* 2017) promises to advance these questions. In the end, this might enable moving beyond 'one-size-fits-all' brain stimulation by using more targeted tDCS protocols (e.g. high-definition tDCS; Datta *et al.* 2008; Dmochowski *et al.* 2011; Edwards *et al.* 2013; Kuo *et al.* 2013), informed by individuals' anatomy, and based on a better understanding of the dose-response.

9.4 General conclusions

In summary, the work presented in this thesis shows that: (1) M1 GABAergic inhibition causally determines the time constant of adaptation memory persistence, explaining phenotypic variation in adaptation decay in young healthy adults; (2) age-related decline in M1 GABA explains the increase in adaptation memory persistence in older adults, illustrating the maintenance of this mechanism in the ageing human brain; (3) spatial neglect is associated with structural and functional alterations of the attention networks; and (4) repeated sessions of prism adaptation combined with M1 anodal transcranial direct current stimulation is feasible in spatial neglect patients but its clinical benefit over the standard multi-session prism therapy has yet to be demonstrated.

Appendix A

Bayesian Decision Theory

In this Appendix chapter, we present the Bayesian Decision Theory framework and discuss how it can be applied to prism adaptation data. It is adapted from:

Petitot, P., O'Reilly, J. X. & O'Shea, J. Towards a neuro-computational account of prism adaptation. *Neuropsychologia* (2017)

A.1 Internal models for sensorimotor control

Most theoretical models of sensorimotor control posit that prior to generating a successful action (e.g. pointing to a visual target), the nervous system should produce a coherent, accurate and unbiased estimate of the current state of one's body, the external world, and how they interact, based on all useful sources of information available. Based on this knowledge, an action plan can then be selected, one that is most likely to maximise performance in relation to the current behavioural goal. The selection of the most appropriate action plan is fundamentally a decision process that is tightly coupled to the ability to generate accurate predictions about the consequences of one's actions on the world and/or one's body (Körding & Wolpert 2006).

Within this framework, sensorimotor control is proposed to rely on internal models

of the external world and the mechanics of the body. Inverse motor models transform a desired behavioural goal into an action plan to accomplish it (Wolpert & Kawato 1998; Kawato 1999). When applied to the sensory domain, inverse models infer the current state of the body (e.g. position of the hand in space) from sensory input (e.g. proprioceptive feedback). Forward models work in the opposite direction and predict the next state (e.g. next position of the hand) based on an estimate of the current state, a copy of the motor command (efference copy), and some internal representation of the complex causal relationship between the two (for review, see Miall & Wolpert 1996; Davidson & Wolpert 2005; Lalazar & Vaadia 2008; Shadmehr *et al.* 2010; Franklin & Wolpert 2011). In the context of this thesis, we will distinguish between a visual and a proprioceptive forward model, each of which generates predictions about the likely next (i.e. expected) visual and proprioceptive state. One advantage of continuously predicting the next state of the body is that this limits the impact of neural transmission delays inherent in relying on actual sensory feedback instead of feedforward predictions (Shadmehr *et al.* 2010). It also allows the brain to compare continuously the veridicality of its predictions against the actually observed (i.e. sensed) measure of a state (provided by sensory inverse models). Deviations between predicted and observed states (i.e. prediction errors) can be used as a signal to drive updating of internal models.

In this Appendix, we outline a theoretical framework that offers a mathematical description of the concepts introduced above. Bayesian decision theory is composed of two components, Bayesian statistics and decision theory. Bayesian statistics offers a means to formalise how an ideal observer should combine new information (e.g. observed sensory input) with prior beliefs (e.g. predicted sensory input), and how multiple sources of uncertain information (e.g. multiple sensory modalities, predictions, prior knowledge) should be integrated, in an *optimal* fashion, to generate a

more certain combined estimate of the current state. Decision theory describes the process of rationally selecting actions, based on the predictions of internal models, given the current behavioural goal. Bayesian decision theory therefore provides a general framework to formalise optimal state estimation and action selection in a dynamic and uncertain world.

A.2 Decision theory and rational movement planning

Any given action (e.g. grasping a cup, hitting a ball) can be executed by an almost infinite number of possible movements ([Figure A.1](#)). Yet, individuals often move in a very stereotypical way (Morasso 1981; Wolpert & Kawato 1998; Simmons & Demiris 2005). Why is this so? Decision theory provides a mathematical framework that formally explains why individuals 'choose' to move the way they do. The central concept is that the cost of a potential movement (e.g. energy consumed, fatigue, risk of an injury, etc.) is weighed against the potential reward that is expected from that movement (e.g. sporting success, monetary gain, altruistic feelings, etc.) (Mazzoni *et al.* 2007; Shadmehr *et al.* 2016). Mathematically, the concept of a utility function (or cost function, the negative counterpart) captures the complex relationship that integrates all of these factors, and quantifies the overall desirability of a potential movement choice. Within this decision theoretic framework, the process of selecting one action amongst alternatives is operationalised as the rational choice of whichever movement plan that maximises expected utility (or minimises expected cost) (Körding & Wolpert 2006; Berniker & Körding 2011). Mathematically, expected utility is

defined as:

$$\text{Expected Utility}(\text{movement plan}) = p(\text{outcome} \mid \text{movement plan}) \cdot U(\text{outcome}) \quad (\text{A.1})$$

where $p(\text{outcome} \mid \text{movement plan})$ is the current estimate of the probability of obtaining an outcome given a particular movement plan (e.g. probability of reaching the target given a certain aiming direction) and $U(\text{outcome})$ is the utility associated with this outcome. This mathematical definition becomes intuitive if you consider, for example, a gambling task in which individuals have to choose between alternative options that are associated with varying reward probabilities and reward magnitudes (for example, see Hsu *et al.* 2005; Behrens *et al.* 2007b). In this scenario, the expected utility of each action alternative is the probability of that action yielding a reward, multiplied by the reward magnitude. Choosing the option that maximises expected utility is the definition of choosing rationally. This mathematical framework has been shown to provide a good quantitative fit to human behavioural data in this kind of decision making task (Hsu *et al.* 2005; Behrens *et al.* 2007b) and has been used to identify functional brain imaging signals that co-vary with the computation of these decision variables (O'Doherty *et al.* 2003; Daw *et al.* 2006; Hampton *et al.* 2006).

When planning a movement, more than one decision needs to be made (Figure A.1). For example, reaching to visual target requires selection of an aiming location (under prism displacement this need not overlap with the estimated target position see Figure A.1B), as well as movement speed and movement trajectory. These different aspects of movement planning are likely to involve different utility functions, and decision theory offers a quantitative framework within which these can be integrated to determine rational movement decisions. Whereas in the gambling task example, the experimenter imposes a utility function, by setting the reward structure of

the task, in real-life motor control contexts, it is hypothesised that an agent’s rational choice of movements is guided by internal utility functions. Assuming that individuals naturally choose actions that maximise expected utility, then investigating the rules that dictate individuals’ movement properties offers a window onto the internal utility functions that guide action choices. When generating target directed movements, for example, it has been shown that healthy individuals favour the precision rather than smoothness of movement (Harris & Wolpert 1998), so as to adopt a trajectory that minimises what resembles the terminal squared error (Körding & Wolpert 2004b). Movement speed seems to be determined by both a speed-accuracy trade-off and an implicit cost assigned to the metabolic energy consumed to produce the movement (Mazzoni *et al.* 2007; Shadmehr *et al.* 2016). Within a decision theoretic framework, factors that change expected utility should result in a change of movement plan. The prism adaptation task can be conceptualised as a manipulation that changes the expected utility of pointing movements. Based on equation 1, there are two ways of modifying the expected utility: 1) by changing the utility associated with certain movement outcomes ($U(outcome)$), and/or 2) by changing the probability of obtaining a certain outcome given a movement plan ($p(outcome | movement\ plan)$). These two could be argued to map on to the traditional dual-process model distinction between strategic control and spatial realignment, respectively.




Figure A.1: Decision theory. **A.** Any action can be executed in an almost infinite number of ways. For example, many different movement trajectories would bring one's finger from point A to a point B. Decision theory provides a mathematical framework that describes how a rational decision maker should choose among alternative movement parameters based on their relative level of expected utility. **B.** Prism adaptation can be conceptualised as a manipulation that affects the computation of expected utility associated with the disparity between the visual estimate of the position of the target and the multimodal estimate of the hand location at the reach endpoint. Before prism onset, the movement plan with highest expected utility is the one that minimises this disparity (i.e. peak utility centred on zero on trial 1). After prism onset however, the visuo-proprioceptive conflict introduced by the prisms induces a performance error: the (experienced) utility of the executed movement plan does not match the predicted utility. This should result in a shift of utility in the direction opposite to the error when planning the pointing movement on trial 2: now, the movement plan with highest expected utility is one that results in a negative disparity between the visual estimate of the target position and the multimodal estimate of the hand location at the reach endpoint (i.e. the multimodal estimate of the hand location is to the left of the visual estimate of the target).

Strategic control can be re-conceptualised as changing $U(outcome)$, where the outcome is defined as the disparity between the estimate of the hand position at the reach endpoint and the estimate of the target position. Normally, peak utility is when this disparity is zero, i.e. the pointing finger lands on the same location as the (visual) target. However, when an individual corrects strategically for a rightward

displacement, by choosing not to aim at the perceived (right-shifted) target location, but instead to aim left of that perceived location, s/he is effectively defining a new utility function with its peak at a non-zero disparity (Figure A.1B). Evidence that movement plans are sensitive to changes in $U(\text{outcome})$ has been found in an experiment that imposed new utility functions by modulating the monetary value associated with certain movement outcomes (Trommershäuser *et al.* 2003). Under these imposed task constraints, healthy individuals were able to rationally select movement endpoints that maximised the potential reward.

Spatial realignment can be conceptualised, within this framework, as minimisation of the prediction error associated with the expected $p(\text{outcome} \mid \text{movement plan})$. That is, the prism manipulation changes the probability that a planned pointing movement (with a given aiming location) will accurately hit the target ($p(\text{hitting the target} \mid \text{aiming location})$). Because of the optical shift, during the early phase of prism exposure the experienced utility of pointing movements does not match predictions ($U(\text{missing the target by } 10^\circ \text{ to the right}) < U(\text{hitting the target})$), i.e. there is a prediction error in terms of utility. We propose that this signal should induce changes in internal models in order to update the computation of $p(\text{outcome} \mid \text{aiming location})$. This would gradually displace the peak of expected utility towards a different movement plan (i.e. more left-oriented) to achieve the desired (unchanged) outcome (i.e. point accurately at the target).

It is worth noticing that this explanatory framework can thus incorporate the key feature of the traditional theoretical framework, i.e. the dissociation between two alternative ways to reduce motor errors during prism exposure (Redding *et al.* 2005; Redding & Wallace 2006b) (see chapter 1). A potential criticism of this proposed decision theory framework is that it merely re-describes the traditional dual-process

model and does not add anything new to the understanding of prism adaptation. The answer to this objection is threefold. First, there is value in precise quantitative (mathematical) description of algorithms that explain behaviour, by contrast with qualitative description. Re-specifying behaviour at an algorithmic level also offers a means to generate and test hypotheses about the neural implementation of those computations (O’doherly *et al.* 2007; Krakauer *et al.* 2017). Second, the model we propose offers a formal mathematical description of how 'strategic control' and 'spatial realignment' processes interact, which is typically vague or lacking in the traditional framework. Third, this framework incorporates the contribution of both task errors (i.e. breaches of expectancy in terms of the utility of a movement) and reinforcement learning (Sutton & Barto 1998) in driving behavioural change during adaptation (Huberdeau *et al.* 2015; Krakauer *et al.* 2017). A recent adaptation experiment using visuomotor rotation demonstrated this interaction, by showing that rates of adaptation and retention were differentially modulated by monetary rewards and punishments associated with movement outcomes (Galea *et al.* 2015). Changes in expected utility differentially affected the rate of error correction and the magnitude of subsequent after-effects. This computational framework may also be useful in understanding prism adaptation impairments observed in certain neuropsychological disorders. For example, it has been hypothesised that cognitive impairments observed in patients with Parkinson’s disease arise owing to altered utility functions because of depleted dopamine in the basal ganglia (Frank *et al.* 2004; Mazzoni *et al.* 2007; Pekny *et al.* 2015). Prism adaptation impairments, including altered error reduction and reduced after-effects, have also been reported in these patients (Weiner *et al.* 1983; Stern *et al.* 1988; Canavan *et al.* 1990; Fernandez-Ruiz *et al.* 2003). We propose that decision theory offers a unified conceptual framework within which to consider and identify possibly common computations (e.g. altered utility functions) that may explain both the classical cognitive impairments observed in Parkinson’s disease (e.g. apathy) and

concomitantly altered mechanisms of prism adaptation behaviour. Thus, re-casting accounts of prism adaptation behaviour within this computational framework offers a means of advancing conceptual and empirical understanding of cognitive impairments across different paradigms of enquiry.

How is the prediction $p(\text{outcome} \mid \text{movement plan})$ generated, and how is the predictive model (motor, proprioceptive, visual) adjusted during prism adaptation? The following two sections offer a Bayesian answer to these questions.

A.3 Bayesian statistics and optimal state estimation

In order to generate appropriate motor commands given a movement plan, the nervous system needs to generate, update, and maintain accurate estimates of states of the body and the external world as we move through space (Franklin & Wolpert 2011). During a pointing movement, multiple sources of error feedback (visual, proprioceptive, motor), with different time constants, both in-flight and following the reach endpoint, provide information about the hand and target position, and the discrepancy between the two. In addition, forward models generate advance predictions about expected outcomes before sensory observations occur. The nervous system must combine these multiple sources of information (predicted and observed) in order to generate a single integrated state estimate (e.g. position of the hand in space) that will guide behaviour (for review, see Ernst & Bühlhoff 2004; Chandrasekaran 2017). The uncertainty associated with any source of information places the problem of state estimation within a statistical framework. Bayesian statistics posit that inference about the state of the body or the external world can be made by combining and weighting each source of information according to its relative level of reliability. Consider for example the task of estimating the reach endpoint position of

the hand after executing a pointing movement. Sensory predictions generated by forward models (e.g. proprioceptive forward model) constitute a prior assumption about where the hand should be if our internal models are correct. In Bayesian terms, such prior knowledge is represented as a prior distribution (blue curve in [Figure A.2A](#)). The sensory system also provides (noisy) information about the likely position of the hand (e.g. proprioceptive feedback). This information is represented as a likelihood distribution (dotted black curve in [Figure A.2A](#)). Combining both sources of information (prior and likelihood) according to Bayes' rule offers a way to calculate a single, more certain estimate of the hand position, represented as a posterior distribution (red curve in [Figure A.2A](#)). Because it minimises uncertainty, such estimate is called optimal. Bayes' rule states that the probability of the hand being in a location x , given a proprioceptive observation o , is the product of the prior probability of the hand location (i.e. prediction about the most likely position of the hand) and the likelihood (i.e. visual estimate of the hand location), normalised by the probability of the observation. Mathematically, Bayes' rule defines the posterior as follows:

$$p(x|o) = \frac{p(x) \cdot p(o|x)}{p(o)} \quad (\text{A.2})$$


Several studies have shown that, when integrating multiple sources of uncertain information, humans behave in a near optimal way that is well predicted by Bayesian models (Ernst & Banks 2002; Alais & Burr 2004; Körding & Wolpert 2004a; Sober & Sabes 2005; Körding *et al.* 2007; Kayser & Shams 2015). [Figure A.2B](#) illustrates the use of Bayes' rule for visual-proprioceptive integration under the assumption of normally distributed noise. The uncertainty (i.e. width of the Gaussian curve) associated with the proprioceptive estimate has been found to be greater than that of its visual homologue (van Beers *et al.* 1996; van Beers *et al.* 1999; Ernst & Banks 2002; Burge *et al.* 2010). Thus, according to Bayes' rule, the optimal combination of pro-

proprioceptive and visual estimates of the hand location should result in a multimodal estimate that is biased towards the visual estimate.


A) Bayesian integration of a forward prediction with a sensory observation

Case 1: no prediction (passive movement)


Case 2: prediction (active movement)



B) Without prisms, passive condition



C) With prisms, passive condition



D) With prisms, active condition




Figure A.2: Bayesian statistics in sensorimotor control. Bayesian statistics describe how multiple sources of uncertain information can be combined optimally into a joint estimate. Here we consider the example of estimating the location of the hand in space. For all panels, the x-axis represents all possible locations of the hand and the y-axis is the probability of the hand being positioned at this location. **A.** The optimal integration of a forward prediction (prior, in blue) with a sensory observation (likelihood, in black) is illustrated under two conditions. During a passive movement (case 1), no forward prediction is generated, i.e. the prior distribution is flat. The resulting sensory posterior estimate (in red) will therefore have the same distribution as the observation (likelihood). During an active movement however, a forward model generates a prediction of the next most likely hand position (prior) that can be integrated with sensory feedback (likelihood) to refine state estimation. This results in a posterior estimate (in red) that is more certain than the likelihood or the prior. **B-D** Bayesian statistics can also be used to describe (multimodal) visual-proprioceptive integration as illustrated here. The black dot on the x-axis represents the true location of the hand. **B.** In the absence of prisms, the visual (green) and proprioceptive (red) estimates of the hand location are close to each other. Based on these two estimates, Bayes' rule can be used to compute a posterior estimate of the hand location (purple) that is more certain than any of the sensory estimates. It posits that the relative level of uncertainty of the two sources of information (visual estimate and proprioceptive estimate) determines their relative contribution to the posterior estimate. Because the visual domain is more reliable (i.e. the width of the distribution is narrower), the resulting multimodal posterior estimate is biased towards the visual estimate. **C.** The visuo-proprioceptive conflict induced by prism glasses is illustrated as a shift of the visual estimate (green) in the direction of the prismatic shift (towards the right). Bayes' rule predicts that the resulting multimodal posterior estimate of the hand location is biased towards the visual observation because it is more reliable (i.e. lower standard deviation of the estimate) than the proprioceptive observation; this effect is called visual capture. The contribution of visual versus proprioceptive estimates of hand location to the multimodal posterior estimate is determined by their relative level of uncertainty. **D.** In this condition, individuals actively move their limb without visual feedback prior to locating their hand in space. The execution of an active movement without visual feedback generates a proprioceptive forward prediction (case 2 of panel A) but no visual forward prediction (case 1 of panel A). If the proprioceptive forward prediction is accurate, the confidence in the resulting estimate will be increased. According to Bayes' rule, the proprioceptive estimate will therefore have a greater influence on the multimodal posterior estimate of the hand location. In other words, the magnitude of visual capture is reduced.

In the context of prism adaptation, Bayesian statistics thus allow a re-formulation in mathematical terms of the perceptual phenomenon known as ‘visual capture’. Visual capture refers to the finding that individuals report their limb to be located closer to where it looks than where it feels under the visuo-proprioceptive conflict generated by prisms (Hay & Pick Jr 1966; van Beers *et al.* 1999). Traditionally, this phenomenon has been described qualitatively as “a kind of perceptual fusion of the discrepant visual and proprioceptive stimuli” where “the proprioceptive stimulus [...] fail[s] to evoke its normal response” (Hay & Pick Jr 1966). By providing a rule to

define optimal multisensory integration, Bayes' theory makes quantitative predictions about the expected magnitude of sensory capture based on the level of uncertainty associated with information provided by each sensory modality (van Beers *et al.* 1999; Ernst & Banks 2002; Burge *et al.* 2010) (Figure A.2BC). It also offers a theoretical explanation for the reduction in magnitude of visual capture that has been observed following an active hand movement compared to a passive one (Welch *et al.* 1979). The execution of a motor command elicits visual and proprioceptive forward models (efference copy) to predict the most likely next visual and proprioceptive observations (Lalazar & Vaadia 2008). These predictions provide additional sources of information that can be combined with actual sensory observations to refine state estimation, i.e. reduce the uncertainty about the resulting percept. Because no visual feedback was provided during movement execution in Welch and colleagues' experiment (Welch *et al.* 1979), only the proprioceptive estimate of the hand position could benefit from integration of the prediction generated by a (proprioceptive) forward model. We suggest that this explains why the resulting percept was more biased towards the proprioceptive estimate (i.e. lower visual capture) in this condition (Figure A.2D).

Visual capture has traditionally been studied relatively independently of prism adaptation, as it does not directly produce after-effects (Held & Hein 1958; Welch *et al.* 1979). The next section, however, will illustrate how, when expressed in Bayesian terms, the phenomenon of 'visual capture' is an important factor that should influence the relative extent of proprioceptive versus visual internal model updating that occurs during prism adaptation.

A.4 Internal model updating and credit assignment

Prism after-effects occur only when active but not passive movements are executed under prism exposure (Held & Hein 1958; Welch *et al.* 1979). This is consistent with the idea that the generation of predictions prior to movement execution is fundamental to the ability to adapt in a way that generates after-effects. In a situation in which the reliability of sensory information is believed to be unchanged, persistent differences between predicted and measured observations implies the need for an update of internal models (Lalazar & Vaadia 2008). The goal of this updating is to minimise systematic prediction errors and maintain internal consistency between predictions and observations.

In order to determine which internal model(s) to update (motor, visual, proprioceptive), the rate at which it should be updated, and for how long the update should be retained, the brain needs to solve what is known as a credit assignment problem, i.e. estimate the underlying cause of the prediction error, in order to assign learning to the correct internal model and update it appropriately (Wolpert *et al.* 2011). It has been proposed that the way in which credit is assigned can be inferred behaviourally by predicted effects on: the rate of error reduction, or the pattern of generalisation (to other contexts, spatial locations, body parts, modalities), or the timescale over which adaptation memory is retained (Cothros *et al.* 2006; Körding *et al.* 2007; Berniker & Körding 2008; Kluzik *et al.* 2008; Diedrichsen & White 2010).

When pointing towards visual targets while wearing prism glasses, there are many possible perturbations that could equally explain the occurrence of prediction errors. For example, a change in the arm dynamics – caused by fatigue for example – can alter the implementation of the motor command, causing systematic deviations from the predicted state. In this case, the nervous system should modify the relationship

between a desired state and the associated motor command, i.e. adapt a motor internal model. Alternatively, prediction errors can arise from incorrect sensory estimates of the target location and/or of the position of the hand. In such a situation, the motor internal model does not necessarily need to be updated, but the sensory (visual and/or proprioceptive) internal models do.

Credit assignment refers to the complex problem of attributing an error to its causal source (Kluzik *et al.* 2008). It requires the nervous system to take into account a large number of parameters, which could include contextual cues, volatility of the environment, and uncertainty in the estimates, predictions, and sensory observations. Computational models incorporating a Bayesian estimator of the source of the error signal have been able to successfully explain post-adaptation generalisation (Berniker & Kording 2008; Haith *et al.* 2009). For example, using a Bayesian estimator of the visual, proprioceptive and motor contributions to the overall perturbation, Haith *et al.* (2009) generated and confirmed the un-intuitive prediction that even a purely motor disturbance, involving no intersensory conflict (force-field), would lead to sensory adaptation, because of uncertainty in the source estimates. Thus, to determine which internal model to update (visual, proprioceptive, or motor), it appears crucial to have a computational architecture in which the uncertainty associated with each source of information is represented (Berniker & Kording 2011). [Figure A.3](#) illustrates such architecture. Under most normal daily circumstances (e.g. not wearing prisms), the uncertainty associated with the predicted landing position of the finger (output of level 3 in [Figure A.3](#)) makes the motor system the most likely source of the error, which should cause preferential updating of the motor inverse model (Kawato 1999; Körding *et al.* 2007; Inoue *et al.* 2016). If credit assignment proceeds in a statistical fashion, a non-zero probability that the error arose from faulty sensory internal models should be inferred. In order to determine which of the visual or proprio-

ceptive internal models is more likely to be responsible for the prediction error, the modality-specific and cross-modal levels of integration (levels 7 and 8 in [Figure A.3](#) respectively) should be considered. At the modality-specific levels of integration (i.e. visual prior/likelihood, proprioceptive prior/likelihood, see levels 7a and 7b of [Figure A.3](#)), information theory offers a way to determine whether the prediction error (i.e. divergence between prior and likelihood) should be attributed to faulty forward or inverse models. Low confidence in the prediction (i.e. uncertain prior) and/or high confidence in the sensory observation (i.e. certain likelihood) imply the need for a sensory forward model update. Alternatively, high confidence in the prediction (i.e. certain prior) and/or low confidence in the sensory observation (i.e. uncertain likelihood) indicate the need for a sensory inverse model update. Degrading the quality of the visual feedback of the reach endpoint error (e.g. by blurring the target or by reducing the duration of the terminal visual feedback) is a way to experimentally manipulate the level of reliability of the visual observations, hence thus potentially maximising sensory inverse model updating. At the cross-modal level of integration (level 8 of [Figure A.3](#)), comparison of the uncertainty in the visual versus proprioceptive estimates should guide the neural system towards an update of the visual proprioceptive model. In the case of a visuo-proprioceptive conflict (i.e. divergence between the visual and proprioceptive estimates), the sensory internal models associated with the most uncertainty (visual or proprioceptive) are the ones that should be preferentially updated. Because this computational architecture ([Figure A.3](#)) incorporates the level of uncertainty associated with each source of information (width of the Gaussian curve) at every level of integration, there is sufficient information to determine which internal model (visual, proprioceptive and/or motor) should be updated, and how it should be updated.




Figure A.3: A Bayesian account of prism adaptation. This schema illustrates putative information processing occurring during the first reaching movement (closed-loop pointing) on trial 1 of prism exposure. First, light from the target is refracted through the prism lens and enters the eye inducing sensory input (yellow arrow between the target and the eye, 1). Based on this input, a visual inverse model generates an estimate of the most likely location of the target (yellow arrow between the visual inverse model and the visual estimate of the target location, top of figure, 2). Owing to the prismatic displacement, this estimate will be right-shifted relative to the true location of the target in space (represented as the red dot, which is left of the visual estimate of the target location). The manual aiming direction believed to maximise expected utility, i.e. that which is expected to successfully align the reach endpoint with the target location, is selected by a decision maker (3). This process is explained in detail in [Figure A.1](#). The selected aiming location is then fed to a motor inverse model, which transforms the desired goal (aiming direction) into an action plan to accomplish it (motor command, 4). An efference copy of this motor command is sent to visual (5a) and proprioceptive (5b) forward models that generate modality-specific predictions about the most likely next location of the hand. In this example, because a ballistic movement is generated (closed-loop pointing trial), this prediction concerns the hand position at the reach endpoint. Meanwhile, the execution of the motor command by the muscles generates visual and proprioceptive feedback (depicted by the arrows originating from the muscle spindle and eye symbols) that is integrated by sensory inverse models to generate modality-specific estimates of the sensed location of the hand (6a: visual and 6b: proprioceptive estimate of the hand location at the reach endpoint). Bayesian integration of the modality-specific forward predictions (prior) with the sensory observations (likelihood) generates modality-specific estimates of the most likely hand location (posterior) (7a and 7b). See [Figure A.2A](#) for a detailed description of this type of integration. The resulting visual and proprioceptive (posterior) estimates are then combined into a multimodal (visuo-proprioceptive) posterior estimate of the hand position (8). See [Figure A.2B-D](#) for a detailed description of this type of integration. On trial 1 of prism exposure, the rightward prism displacement induces: 1) leftward performance error (i.e. divergence between the visual estimate of the target position, 2, and the multimodal estimate of the hand location after movement execution, 8); 2) prediction errors in the sensory internal models (i.e. divergence between prior and likelihood, 7a-b); 3) visuo-proprioceptive conflict (i.e. divergence between the proprioceptive and visual estimates of the hand location, 8). The effect of internal model updating is to reduce these performance errors and to gradually realign all the distributions depicted in this schema (i.e. aiming location, predicted visual and proprioceptive observations, visual and proprioceptive observations, visual and proprioceptive estimates, multimodal estimate). Internal models are updated iteratively within this network as evidence accumulates and performance is adjusted from one trial to the next. In order to update internal models appropriately, the brain has to attribute learning to the correct internal models (motor, visual or proprioceptive). The level of uncertainty associated with each source of information (sensory predictions, sensory observations, sensory estimates) at every level of integration (modality-specific and multimodal) is used to infer the most likely error source and hence assign learning to the correct internal model. In this example, the visual estimate is more certain than the proprioceptive estimate, which should result in a preferential update of proprioceptive internal models. This model architecture incorporates the ‘strategic control versus spatial realignment’ dissociation as performance error can alternatively be quickly corrected at the level of the decision maker by selecting an aiming location that would result in a negative disparity between the visual estimate of the target position and the multimodal estimate of the reach endpoint hand location (see [Figure A.1](#) for a detailed description).

Compared to the traditional dual-process model, this framework makes quantitative predictions and offers a principled explanation of the predicted behavioural consequences of modulating the uncertainty of specific information sources (for example, see K rding & Wolpert 2004a; Burge *et al.* 2010; Yamamoto & Ando 2012). For example, this can explain why terminal prism exposure (no visual feedback during reaching movements, endpoint error feedback only) induces greater visual than proprioceptive after-effects, while concurrent prism exposure (visual feedback during reach movements and endpoint error) induces greater proprioceptive than visual after-effects (Redding & Wallace 1996; Redding & Wallace 2001; Redding & Wallace 2006b). The occlusion of visual feedback under terminal prism exposure conditions increases the uncertainty of the visual estimate of the hand position. Hence, when the prisms induce a prediction error, the cause of this is more likely to be attributed to a faulty visual internal model, thus resulting in greater visual than proprioceptive model updating. By contrast, the presence of visual feedback in concurrent prism exposure conditions allows for continuous state estimation to occur in both the visual and proprioceptive domains. It is known in such condition that visual estimates are more certain relative to proprioceptive estimates (Figure A.2; Ernst & Banks 2002; Burge *et al.* 2010). Hence, when a prediction error is experienced at the reach endpoint (under concurrent prism exposure), it is more likely to be attributed to a faulty proprioceptive internal model, which is then updated accordingly. Based on this interpretative framework, we predict that, across individuals, the magnitude of visual capture (which quantifies relative confidence in visual versus proprioceptive information) should predict the degree to which individuals update their visual relative to their proprioceptive internal models under concurrent prism exposure (i.e. the relative difference in after-effect magnitude between visual straight ahead judgement and proprioceptive straight ahead pointing).

Appendix B

Testing the inter-hemispheric competition account of visuospatial attention with combined TMS/fMRI

This chapter evaluated the role of the right posterior parietal cortex in distributing visual attention across space. Combined rTMS and fMRI was used to assess the behavioural and physiological consequences of an inhibitory stimulation of the right posterior parietal cortex. We provided evidence for a rightward shift in visuospatial attention accompanied by a leftward shift in inter-hemispheric dominance. These results support the inter-hemispheric competition account of visual extinction. The core of this chapter is adapted from:

Petit, P., Noonan, M. P., Bridge, H., O'Reilly, J. X. & O'Shea, J. Testing the inter-hemispheric competition account of visual extinction with combined TMS/fMRI. *Neuropsychologia* **74**, 63–73 (2015)

B.1 Introduction

Following injury to right parietal cortex, patients often exhibit visual attentional dysfunction, such as neglect of stimuli in left space. In a related syndrome, extinction, processing of left stimuli is intact, but if presented simultaneously with a stimulus on the right, the stimulus on the right is detected while that on the left is ‘extinguished’ from awareness. Theoretical models of the pathophysiology of neglect and extinction emphasise the critical role of relative hemispheric dominance and inter-hemispheric competition in controlling the allocation of attention across space (Kinsbourne 1976; Heilman & Van Den Abell 1980; Mesulam 1981; Cohen *et al.* 1994). The inter-hemispheric competition model posits that the left and right parietal cortices compete to direct attention towards contralateral hemispace, with each hemisphere exerting an inhibitory influence over the other. In a healthy brain, these competing attentional vectors are thought to be broadly counterbalanced, enabling attention to be distributed across both hemifields. However, correlative evidence from brain imaging in healthy volunteers suggests that the right inferior parietal lobe is dominant over the left during bilateral attention (çiçek *et al.* 2007), consistent with the greater severity of attentional impairments after right than left injury (Weintraub & Mesulam 1987). Right parietal damage both directly weakens leftward attention, and indirectly, via transcallosal disinhibition, leads to hyper-activation of left parietal cortex, consequently exacerbating a rightward attentional bias. Thus, disrupted parietal inter-hemispheric balance is thought to deviate attention rightward, resulting in a competitive advantage for right hemifield stimuli during tasks that require bilateral attention (de Haan *et al.* 2012).

Brain imaging and brain stimulation studies have confirmed a role for disrupted inter-parietal hemispheric balance in contributing to neglect/extinction in the damaged brain (Corbetta & Shulman 2002; Corbetta *et al.* 2005; Koch *et al.* 2008). How-

ever, these models of pathology also make claims about normal function: specifically, that it is the relative balance of activity between left and right parietal cortex that causally mediates bilateral spatial attention. Brain imaging studies have provided confirmatory correlative evidence that participants' behavioural bias between left and right space relates to measures of functional lateralisation within attention structures (De Schotten *et al.* 2011; Szczepanski & Kastner 2013; Benwell *et al.* 2014). Causal evidence has come from studies using non-invasive transcranial magnetic stimulation (TMS) to transiently perturb function in parietal cortex. Based on the observed pattern of TMS-induced behavioural interference, inferences have been made about the presumed nature of functional interactions between left and right parietal cortex.

The first two TMS studies of this kind used a visual detection task, designed to mimic clinical confrontation testing of extinction, combined with online high-frequency (Pascual-Leone *et al.* 1994) or offline low-frequency rTMS (Hilgetag *et al.* 2001). Participants had to monitor both hemifields simultaneously to detect upcoming transient targets, which appeared either on the left, right or bilaterally. In both studies, right parietal TMS impaired detection performance on bilateral trials. In Hilgetag *et al.*'s study, behaviour was assessed before and after 10 minutes of 1Hz TMS applied at rest. The key finding was that right parietal TMS induced a rightward shift in the spatial distribution of participants' errors on bilateral trials. That is, on trials in which participants failed to detect both targets, after TMS they were more likely to detect only the target on the right and omit that on the left, similar to clinical extinction. Subsequent behavioural TMS studies have replicated the finding that right parietal interference impairs detection of left stimuli on bilateral trials (Dambeck *et al.* 2006; Meister *et al.* 2006; Duecker & Sack 2015). However, the functional basis of impaired performance on bilateral trials is unclear, since an attention deficit for left space would itself compromise bilateral trials, which are the

most attentionally demanding and therefore the most sensitive trial type. Hence, existing TMS evidence for a specific causal role of parietal inter-hemispheric balance in mediating bilateral attention currently relies on reverse inference from behavioural impairments and is therefore weak.

Rather than inferring function from behaviour, we combined TMS with fMRI to test directly the hypothesis that normal bilateral attentional function depends on the balance of activity between left and right parietal cortex. We used a visual target detection task previously used to simulate extinction using TMS in the healthy brain (Hilgetag *et al.* 2001). The aim was to use a task in the MRI scanner that would provide a sensitive means of detecting changes in the balance of parietal inter-hemispheric competition caused by TMS, even if any behavioural effects are likely to be subtle or short-lived.

We targeted stimulation at the caudal part of the angular gyrus (ANG) at the junction with the intra-parietal sulcus (IPS). We therefore refer to the TMS protocol as targeting right ANG/IPS. Stimulation was applied to the right hemisphere, given the dominance of right parietal cortex over left in extinction and neglect (Becker & Karnath 2007; çiçek *et al.* 2007), and since previous TMS studies have implicated this region in attention (Ashbridge *et al.* 1997; Rushworth *et al.* 2001) and shown that stimulation at this site can induce extinction-like behaviour in healthy individuals.

First, we functionally localised a region of right posterior parietal cortex in each individual at which online high frequency repetitive TMS (10Hz, 500ms) disrupted attentional performance (Ashbridge *et al.* 1997). Anatomical mapping confirmed that this 'hotspot' was located in the caudal part of the right ANG at the junction with the IPS. In a subset of participants, we then confirmed that offline low frequency

TMS (1Hz, 15 minutes) to that brain region induced extinction-like behaviour in an orthogonal visual detection task. Whereas TMS did not change detection accuracy, the spatial distribution of errors on bilateral trials was shifted from left to right, simulating clinical visual extinction and replicating the key finding of Hilgetag et al. (2001). We then conducted the main fMRI experiment, in which participants performed the same visual detection task, once before and after 1Hz TMS was applied to transiently inhibit function of the right ANG/IPS. Given the behavioural finding that this stimulation protocol shifted attentional bias from left to right, we analysed the functional imaging data to test the hypothesis that right ANG/IPS TMS would shift the parietal inter-hemispheric balance from right to left.

B.2 Material and methods

B.2.1 Participants

Five right-handed individuals participated in the behavioural TMS experiment (4 male, mean age = 23.8 years, SD = 4.7). Twelve right-handed individuals participated in the main TMS/fMRI experiment (4 male, mean age = 25.9 years, SD = 3.5), four of whom also underwent the behavioural TMS experiment. All participants gave written informed consent as approved by the Central Office for Research Ethics Committee (COREC reference number: 05/Q1606/96). All participants had normal or corrected-to-normal vision and indicated no family history of psychiatric or neurological disease. No participant reported any side effect from the experimental procedure.

B.2.2 TMS functional localisation of right posterior parietal cortex

All participants in both experiments first underwent a 2-hour functional localiser session to determine the anatomical target for TMS in each individual. To identify the sub-region of right angular gyrus/intra-parietal sulcus (rANG/IPS) at which stimulation would disrupt attentional function, participants underwent an established mapping protocol previously shown to identify parietal regions functionally involved in the allocation of visual attention (Ashbridge *et al.* 1997; Rushworth *et al.* 2001). That location was then targeted for TMS in the main experiments.

Each trial of the functional localisation procedure started with an alerting tone and a white fixation spot (500ms) at a random location on the screen, followed by the search array (750ms), composed of red and green diagonal lines on a black background (Figure B.1A). The target was a red line oriented at 45 surrounded by two kinds of distracters – green lines oriented at 45 and red lines oriented at 135. The target was present on half of trials. Participants were instructed to respond on each trial with a right index finger button press if they detected the target and a right middle finger button press if not. The inter-trial interval lasted until participants responded or up until 4 seconds maximum. The visual search task was run on a computer with a Windows 98 operating system (75Hz refresh rate 1024 × 768 resolution) triggered by Turbo Pascal (Version 7.0, Borland International, Inc).

Each participant first performed 10 practice task blocks (10 trials per block), the last 5 of which were performed with sham TMS (coil placed on the head but oriented away), in order to habituate participants to the somatosensory and acoustic artefacts of the subsequent TMS procedure. Earplugs were worn throughout. Participants then performed blocks of the search task, during which a train of high frequency

TMS (10Hz, 500ms, 65% maximum stimulator output, 70mm coil) was applied time-locked to search array onset. TMS blocks alternated with baseline blocks, in which no TMS was applied. The goal was to identify the scalp location overlying parietal cortex in each individual at which TMS slowed reaction time ($> 20\text{ms}$ delay relative to the preceding no TMS block). Illustrative data are shown in [Figure B.1C](#). TMS was applied iteratively to each of nine points on a $3 \times 3\text{cm}$ square lattice marked on the scalp and centered over electrode position P4 (ie: 9cm superior from the inion and 6cm lateral). Points were sampled until the attention ‘hotspot’ location was found.

The location of the right ANG/IPS hotspot in each individual was then confirmed anatomically using Brainsight frameless stereotaxy (version 1.5 Rogue Research, Canada). Each participant’s head was co-registered with their anatomical MRI in native space, and a trajectory was plotted from the TMS ‘hotspot’ scalp location onto the cortical surface. Individual participants’ structural MRI scans were then normalised to the MNI 152-mean brain T1 template. In one individual no attention ‘hotspot’ location could be found, so TMS was instead targeted at the group mean anatomical ‘hotspot’ coordinate reported in previous papers using this procedure (ie: MNI coordinates $X = 41, Y = -69, Z = 43$) ([Rushworth *et al.* 2001](#)). In the current study the group mean ($\pm SD$) coordinates at which TMS was applied were: $X = 39.60(\pm 8.85), Y = -61.80(\pm 6.30), Z = 51.80(\pm 6.87)$ in the behavioural TMS experiment ($n = 5$) and $X = 41.67(\pm 6.76), Y = -64.17(\pm 7.87), Z = 48.08(\pm 6.69)$ in the TMS/fMRI experiment ($n = 12$) (see [Figure B.1](#)).




Figure B.1: Functional localisation of parietal cortex and the visual extinction task. **A.** Schematic of the conjunction visual search task (Ashbridge *et al.* 1997). A trial began with an alerting tone and a fixation cross (500ms) at a variable location on the screen, followed by the search array (750ms). Participants had to detect and report the presence or absence of a singleton target (red line oriented at 135° indicated by the dotted circle) amongst distractors of two types: red lines oriented at 45° and green lines oriented at 135° using a button press. The inter-trial interval terminated with the response, or lasted up to a maximum of 4 seconds. **B.** Anatomical location of TMS site. Each circle represents the MNI coordinates identified by the functional localisation procedure for each individual participant site at which TMS was applied to the right ANG/IPS in experiments 1 and 2. For illustration purposes, the y- and z-coordinates for each individual were projected against the group mean x-coordinate. The line weight is proportional to the number of participants for which the stimulation coordinates are shown. The stimulation sites cluster around the posterior part of the angular gyrus and inferior border of the intra-parietal sulcus. Coordinate range: $34 < x < 57$, $-75 < y < -47$, $43 < z < 60$. **C.** TMS to right ANG/IPS delays reaction time. To illustrate typical mean reaction times on the visual search task, graph shows data for the no TMS versus TMS condition from experiment 1, illustrating the >20ms reaction time deficit that was the criterion to identify the right ANG/IPS attention 'hotspot'. **D.** Timeline of a single trial of the extinction task used in both experiments.

(Legend continued on next page)

The onset of the central fixation cue onset signaled to participants to monitor both hemifields to detect and report the location of an upcoming transient visual target, which could appear on the left, right or bilaterally, or not at all. This example shows a bilateral target trial. The inter-trial interval (ITI) varied across experiments. **E.** Staircase procedure to determine stimulus size for the extinction task. To illustrate the psychophysical thresholding procedure, graph shows the mean percentage accuracy for left, bilateral and right target trials across 5 different stimulus sizes (in arbitrary units, a.u.) from experiment 1. Two adjacent stimulus sizes were chosen per participant based on bilateral trial detection accuracy of 40-60%.

B.2.3 Transcranial Magnetic Stimulation

Since the main aim of this study was to image the effects of TMS, we chose an offline low-frequency TMS protocol (1Hz, 15 minutes) to transiently inhibit the functioning of right ANG/IPS, similar to Hilgetag et al. (2001). Such offline stimulation protocols have been shown to inhibit neural activity and perturb function for a period that outlasts the stimulation duration by approximately 5-15 minutes (Chen *et al.* 1997; Boroojerdi *et al.* 2000), thus allowing for induced changes in functional brain activity to be detected by subsequent fMRI (O'Shea *et al.* 2007).

In both the behavioural and TMS/fMRI experiments a biphasic Magstim Super Rapid TMS machine (Magstim Company, Whitland, Wales) was used to deliver 15 minutes of 1Hz offline TMS to the right posterior parietal cortex at 90% of resting motor threshold through a 70mm figure-of-eight coil while participants sat at rest. The coil was held tangential to the skull and oriented at 45° from the mid-sagittal axis over the identified rANG/IPS in each individual. The coil was fixed in place with a metal clamp. To ensure minimal coil movement an experimenter supported both the coil and the participant's head throughout stimulation. Coils were changed after 7.5 minutes to prevent overheating. Coil changeover took less than 1 minute.

Individual resting motor threshold (rMT) was determined at the end of the func-

tional localisation session. The rMT of the right motor cortex was determined for each individual as the minimum intensity of single-pulse TMS to the motor ‘hotspot’ that induced motor evoked potentials (MEPs) of $\approx 50\mu\text{V}$ peak-to-peak amplitude from the first dorsal interosseous muscle of the left hand on an average of 5 out of 10 consecutive trials. TMS to the right ANG/IPS was applied at 90% rMT and 10 additional trials confirmed that this stimulation intensity elicited no MEPs. Electromyographic responses were recorded from silver chloride electrodes in a tendon-belly montage. Responses were amplified, filtered and sampled using a CED 1902 amplifier interconnected with a CED 1401 analogue-to-digital converter (Cambridge Electronic Design Ltd., Cambridge, UK). The signals were recorded using Spike 2 (version 3.21) and Signal (version 2.14) programs. The sampling rate was 5000 samples per second and the signal was band-pass filtered between 10 and 1000 Hz.

B.2.4 Visual Extinction task

In order to test for the behavioural (Experiment 1) and functional (Experiment 2) consequences of offline right ANG/IPS TMS, we used the same target detection task used by Hilgetag et al. (2001) to simulate visual extinction in the healthy brain. On each trial, participants were required to maintain central fixation and covertly monitor both the left and right hemifield to detect and report the location of transiently presented peripheral visual targets using a button press (Figure 1D). Each trial began with a central fixation cross (duration randomised between 2-4 seconds) which functioned as a cue, warning the participant that the target stimulus was about to appear. Immediately upon cue offset, the target was presented for 100ms. A target was either presented in the left visual field (‘unilateral left trials’), the right visual field (‘unilateral right trials’), or to both hemifields simultaneously (‘bilateral trials’). In addition, there were ‘catch’ trials in which no target was presented. All four trial types were equiprobable. Participants were required to respond on every trial by

pressing one of four buttons using the four fingers of their right hand to indicate either ‘left’, ‘right’, ‘bilateral’ or ‘none’. Following target offset, a blank screen filled the inter-trial interval (ITI) for a variable duration (2-4 seconds in behavioural TMS Experiment 1; jittered between 6-8 seconds in the TMS/fMRI Experiment 2). Total task duration was approximately 15 minutes (Experiment 1: 144 trials, Experiment 2: 100 trials).

The targets comprised of small grey squares ($167 \times 167 \times 167$ RGB) presented on a lighter grey background ($170 \times 170 \times 170$ RGB). Following a block of task practice (70 trials, 5 minutes), participants underwent a psychophysical staircase procedure to determine appropriate threshold-level target sizes for the main task. Five stimulus sizes were tested for each individual, with squares of 8-21 pixels in length across 90 trials. During the thresholding procedure, participants were tested only on unilateral left/right and bilateral trials (chance performance = 33%). Two adjacent target sizes were chosen for each participant based on detection accuracy on bilateral trials of 40-60% (Figure B.1). After thresholding, participants underwent a second block of task practice in which catch trials were introduced (50 trials, 5 minutes, chance performance = 25%). In case of a participant’s performance failing to exceed chance, target size was adjusted and the block was repeated until criterion was reached. Participants then performed the main task.

During Experiment 1 (behaviour) the task was run on a PC with a Windows 98 operating system (75Hz refresh rate, 1024×768 resolution) using Presentation software (version 0.52, Psychology Software Tools Inc.). Participants were positioned on a chinrest 44cm from the screen and targets were presented at 16cm horizontally and 5cm vertically from central fixation (ie: at 18° eccentricity). During Experiment 2 (TMS/fMRI) the task was run on a PC with a Windows 98 operating system (60Hz

refresh rate) using Presentation software (version 9.9). Stimuli were back-projected onto a screen in the MRI scanner viewed via a mirror above participants (12.5cm: distance from mirror to eyes) with targets projected at a viewing angle of 18°. Button presses were collected on a parallel port response box.

B.2.5 Experimental Procedures

B.2.5.1 Experiment 1: Behavioural effect of right parietal TMS on extinction task

Owing to the anticipated difficulties controlling the visual environment in the MRI scanner, the goal of this prior behavioural experiment was to confirm that the offline right parietal TMS protocol would induce an extinction-like pattern of detection behaviour under psychophysically controlled conditions. To confirm this, participants performed the visual target detection task before and after TMS. Following task practice, participants were dark-adapted for 10 minutes using a blindfold, after which they underwent psychophysical thresholding to determine the appropriate target sizes, followed by another practice block to confirm the threshold was appropriate. The extinction task was then performed twice, once at baseline (baseline task session) and once after 15 minutes of 1Hz TMS (post-TMS task session). TMS was applied to the right ANG/IPS in each individual using Brainsight frameless stereotaxy based on the coordinates derived from the TMS functional localiser procedure performed previously. As it would be necessary for the TMS/fMRI experiment, stimulation order was counterbalanced across the baseline and post-TMS task sessions (2 participants in baseline/post-TMS order, 3 participants in post-TMS/baseline order). For those participants who performed the post-TMS task session first, 45 minutes elapsed between the end of TMS and the start of the baseline task session. Participants performed 144 trials in each session (36 of each trial type) with an average trial duration of 6.1 seconds.

B.2.5.2 Experiment 2: TMS/fMRI of the extinction task

All stimuli, task events and experimental procedures were identical to Experiment 1, except that since the fMRI design was rapid event-related, the inter-trial interval was lengthened (from 2-4 to 6-8 seconds) to ensure that functional brain activity on trial n could be distinguished effectively from trial $n+1$. Hence, average trial duration was approximately 10.1secs (from 8.1 seconds to 12.1 seconds), resulting in 100 trials in total over 15 minutes. Participants first practiced the task outside the scanner for 5 minutes (70 trials), followed by a staircase procedure to titrate target size inside the scanner, followed by the main extinction task during fMRI acquisition. The order of the baseline and post-TMS scan sessions was counterbalanced (6 participants in baseline/post-TMS order, 6 participants in post-TMS/baseline order), both to minimise practice effects and ensure that any changes in brain activity observed after stimulation would not be contaminated by scan order effects. For those participants who performed the post-TMS scan session first, 45 minutes elapsed between the end of TMS and the start of the baseline scan session. Those participants performed an additional 10 practice trials at the start of the post-TMS scan session, since their first trials might be more error prone. TMS was applied in the console room directly adjacent to the scanner room, and the interval between the end of the TMS train and the start of fMRI data acquisition was approximately 4 minutes.

B.2.6 MR image acquisition

Blood oxygenation-level-dependent (BOLD) fMRI images and T1-weighted anatomical images were acquired on a 3T Siemens Trio MR scanner with a maximum gradient strength of 40 mT.m^{-1} at the Oxford University Centre for Clinical Magnetic Resonance Imaging (OCMR). Participants wore earplugs throughout TMS and scanning. All participants underwent two scan sessions (baseline scan session and post-TMS scan session) in counterbalanced order. During each session, 620 axial echo-planar

volumes (25×6 mm slices, TE = 30 ms, TR = 1500 ms, flip angle = 73° FOV = 192×192 mm², matrix = 192×192 , voxel size = $3 \times 3 \times 5$ mm³) were acquired over 15 minutes. In addition, a 9-minute high resolution structural MRI scan was also obtained (FLASH; repetition time = 3000 ms, echo time = 4.71 ms, and flip angle = 80° giving a voxel size of $1 \times 1 \times 1$ mm³).

B.2.7 Functional magnetic resonance imaging data analysis

All image processing and statistical analyses were performed using tools from the Oxford Centre for Functional Magnetic Resonance Imaging of the Brain Software Library (FMRIB, Oxford, U.K; www.fmrib.ox.ac.uk/fsl). The first two volumes of each fMRI scan were discarded to allow for T1 equilibration effects. The functional MRI data were then corrected for motion using MCFLIRT (Jenkinson *et al.* 2002) and non-brain structures were removed using BET (Smith 2002). The images were spatially smoothed using a 5 mm Gaussian kernel of full-width at half maximum and low-frequency drifts were removed with high-pass temporal filtering with a cut-off period of 100 seconds. Functional data were registered to participants' structural image and MNI152 standard space by using linear transformation (Jenkinson & Smith 2001; Jenkinson *et al.* 2002). fMRI data were denoised using MELODIC (Beckmann & Smith 2004); components that were related to eye blinks and eye movements were removed. The time series data were analysed using a general linear model (GLM) approach. Statistical analysis was carried out in FEAT using FILM with local autocorrelation correction (Woolrich *et al.* 2001). A GLM was performed on a voxelwise basis using a rapid event-related design. The hemodynamic response was modeled as a gamma function, a normalisation of the probability density function of the gamma distribution with zero phase, standard deviation of 3 seconds, and a mean lag of 6 seconds.

The major constraint on the fMRI design was the need to complete scanning within a brief period after offline TMS, in order to maximise the chances of detecting a post-stimulation change in brain activity before it would decay. We opted for a 15-minute maximum scan period, based on our previous work with a similar offline 1Hz 15-minute TMS protocol (O'Shea *et al.* 2007). This necessitated a rapid event-related fMRI design, to maximise the number of different trial types and minimise the amount of rest within this 15-minute period. As a consequence, trial types were short, and it was not possible to disentangle within-trial components. Specifically, it was not possible to dissociate cue-related from target-related activity. Hence, 4 trial type regressors were constructed that reflected a composite of cue- and target-related activity. Each trial type was modeled using a single explanatory variable (EV) that captured the cue period (jittered between 2 and 4 seconds) followed by the target (100 ms) on every trial. This yielded 4 distinct EVs that varied with target location: left, right, bilateral, catch (none). The explanatory variables were modeled with their temporal derivatives. Six additional movement parameters were included to capture the remaining head motion-related variance.

FEAT was used to fit the model to the data, generate parameter estimates for each EV at each voxel, and contrast parameter estimates against the implicit resting baseline. To generate statistical maps of the brain regions specifically engaged in attention and target-related processing (from which motor response-related activity had been subtracted out), we contrasted correct left, right, and bilateral trials versus catch trials.

B.2.8 Statistical analysis of behavioural data

Accuracy and reaction time data from the extinction task were analysed with repeated measures ANOVA and t-tests corrected for multiple comparisons using SPSS

software (SPSS inc. version 22.0). Data distribution assumptions of normality were assessed quantitatively using the Shapiro-Wilk test, and qualitatively using Q-Q plots, while sphericity was assessed using Mauchly's test, with violations corrected using the Huyhn-Feldt procedure where appropriate. Responses to the two target sizes were first calculated independently and then averaged for each individual prior to analysis. To specifically test the prediction that right parietal TMS would induce a rightward shift in attentional orienting, relative detection performance between the left and right hemifields was quantified by a lateralisation index (LI). LI provides a single numerical value that quantifies participants' relative attentional bias between the left and right hemifields. LI was calculated on incorrect responses for bilateral target trials as follow:

$$LI = \frac{(Right - Left)}{(Right + Left)} \quad (B.1)$$

where 'Right' (respectively 'Left') refers to the number of trials for which participants erroneously responded 'Right' (respectively 'Left') instead of 'bilateral'. LI values range between -1 and +1, with 0 indicating no difference between left and right hemifields. Positive LI values indicate a bias towards right space and negative values indicate a bias towards left space. The key prediction was that right parietal TMS would induce an extinction-like behavioural effect on bilateral trials (ie: a relative decrease in the number of incorrect 'left' responses and a concomitant increase in the number of incorrect 'right' responses). That is, we predicted an increase of the bilateral LI (rightward shift) between the baseline and the post-TMS task sessions.

B.3 Results

B.3.1 Behavioural data

B.3.1.1 Experiment 1: Behavioural effect of right parietal TMS on the extinction task

The key prediction was that right ANG/IPS TMS would induce a rightward shift in spatial attention and consequent target detection. Hence, performance on bilateral trials (when participants have to covertly monitor and detect targets in both hemifields) should be the most sensitive measure of any such attentional shift. Specifically, by analogy with clinical extinction, we predicted that TMS would change the spatial distribution of errors on bilateral trials, relatively reducing erroneous 'left' responses while increasing erroneous 'right' responses. To assess this, we tested for a change in the direction of error responses on bilateral trials, as quantified by the lateralisation index (ie: $LI = ('right' - 'left' \text{ incorrect responses}) / ('right' + 'left' \text{ incorrect responses})$), after rANG/IPS TMS compared to baseline.

First we tested for the presence of a lateralised spatial bias at baseline (before TMS was applied). A one-sample t-test (versus zero) indicated that participants were significantly biased towards the left hemifield at baseline, a well-documented phenomenon known as pseudoneglect ($t(4) = 4.518, p = 0.01; mean LI = -0.47$; [Figure B.2C](#)). Next we tested the a priori directional prediction that LI values would shift rightward after TMS (ie: increase). A one-sample t-test (versus zero) on the change in LI bilateral score (delta Post-TMS – Baseline) confirmed that there was a significant rightward shift after TMS ($t(4) = 2.128, p = 0.05, 1\text{-tailed}$; mean δLI bilateral = $+0.503$; [Figure B.2C](#)). The same analyses on unilateral LI scores were not significant (all $p > 0.7$). Hence, right parietal TMS induced a highly specific behavioural effect: it shifted participants' baseline spatial bias from left to right, abolishing pseudo-neglect.

Analysis also considered correct target detection (“hit”) rates for all three trial types in which a target was present (left, right and bilateral). Repeated-measures ANOVA with factors of TMS (baseline versus post-TMS) and Trial Type revealed a main effect of Trial Type ($F(2,8) = 4.54$, $p = 0.048$), reflecting the fact that bilateral target trial performance was lower than unilateral trial performance (Figure B.2). However, there was no effect of TMS ($F(1,4) = 0.31$, $p = 0.61$) and no interaction ($F(2,8) = 0.84$, $p = 0.47$). Analysis of reaction times also confirmed no effect of TMS ($F(1,4) = 0.71$, $p = 0.447$), indicating that the changes observed on bilateral trials were not due to alteration of a speed/accuracy trade-off.

Hence, TMS did not have a generalised non-specific effect on visual sensitivity or response time, but rather had a highly specific effect of changing the spatial distribution of errors on bilateral trials from left to right, thus simulating clinical visual extinction.

B.3.1.2 Experiment 2: TMS/fMRI of the extinction task

The same statistical analyses conducted in Experiment 1 were carried out on the behavioural data in Experiment 2. The overall pattern of detection accuracy in the fMRI experiment was similar to Experiment 1, although baseline hit rates were 9% higher despite the same mean target size being used, indicating the greater difficulty of titrating performance to psychophysical threshold in the scanner environment (Figure B.3A). As in Experiment 1, TMS had no effect on hit rates, reaction times or unilateral LI scores. By contrast with Experiment 1, there was no evidence of baseline pseudoneglect in the scanner ($p > .11$), and consequently no extinction-like effect of TMS ($p > 0.4$).




Figure B.2: Behavioural results of Experiment 1. **A.** No effect of right parietal TMS on extinction task accuracy. Graph shows percentage correct responses across the four trial types of the extinction task performed outside the scanner. There was no difference in accuracy between the baseline (white) and post-TMS sessions. **B.** Graph shows the relative percentage of incorrect responses ('left' 'right', 'none') on bilateral target trials at baseline (white) and after TMS (black). **C.** Graph shows the bilateral Lateralisation Index at baseline (white) and after TMS (black). Prior to stimulation, when participants made an error on bilateral trials, they more frequently detected the target on the left than on the right, ie: they showed pseudoneglect. Analysis showed that TMS abolished this bias, relatively shifting the LI rightward. Asterisk reflects significance level of one sample t-test on the change in LI (delta LI Post TMS – Baseline) against zero (*: $p \leq 0.05$, one-tail).

The absence of a leftward spatial bias ('pseudoneglect') at baseline in the scanner likely reflects the constraints of the visual MRI environment in which accurate psychophysical control over luminance was not possible. Hence, in Experiment 2 it was not possible to investigate neural correlates of, or an effect of TMS on, an absent behavioural effect. However, since behaviour did not differ significantly between the baseline and post-TMS conditions, this meant that any observed changes in fMRI activity could be attributed straightforwardly to a direct causal effect of TMS, without the interpretative complications of concomitant behavioural change.

Hence, whereas Experiment 1 aimed to characterise the causal consequences of TMS for attentional performance under psychophysically controlled visual conditions, the focus of Experiment 2 was physiological, aiming to characterise the causal impact of TMS on functional brain activity mediating attentional performance.

B.3.2 Functional brain imaging data

B.3.2.1 Extinction task network

To identify the network of brain regions activated during bilateral attention, independently of any TMS effect, a multi-subject, multi-session mixed effects analysis was conducted. At the first level, Z statistic images were generated for the contrast [Left/Bilateral/Right - Catch trials], separately for the baseline and the post-TMS scan sessions. Then, a second-level fixed-effects analysis was conducted on each participant's pair of scans to yield a Z statistic map reflecting the average pattern of brain activity across both scan sessions. Finally, a third level mixed-effects analysis was conducted at the group level, on the output from the second level analysis per participant, with Z statistic images thresholded using clusters determined by $Z \geq 2.3$ and a corrected cluster significance threshold of $p = 0.05$. All statistical parametric maps are displayed according to radiological convention (ie: left and right hemispheres are inverted).

Attentional orienting activated a bilateral network of cortical regions including the angular gyrus/intra-parietal sulcus, superior parietal cortex, visual cortex and left precentral gyrus (Figure B.3B, Table B.1). For illustrative purposes, Figure 3 shows the mean location at which TMS was targeted (1cm radius sphere shifted one radius inward from the scalp surface stimulation site; MNI coordinates: $x = 36$, $y = -64$, $z = 42$), demonstrating that stimulation was applied to a region of right ANG/IPS that is functionally recruited when participants are covertly monitoring both hemifields in anticipation of an upcoming target.




Figure B.3: Cortical network activated during bilateral attention (Experiment 2). **A.** No effect of right parietal TMS on extinction task accuracy (experiment 2). Similar to the experiment performed outside the scanner, right parietal TMS had no effect on extinction task performance accuracy during fMRI. **B-C.** Figures show regions activated during the extinction task averaged across both the baseline- and post-TMS scan sessions. Performance of the extinction task activated a bilateral network of regions including the angular gyrus/intra-parietal sulcus (ANG/IPS), superior parietal lobe (SPL) and pre-central gyrus. Activity is time-locked to cue onset, during which participants were required to covertly monitor both hemifields. Voxelwise statistics represent the contrast [Left/Bilateral/Right – Catch trials] averaged across baseline and post-TMS sessions ($Z > 2.3$, corrected cluster extent significance threshold $p < 0.05$). Colour bar represents Z statistics at each voxel. The black circle illustrates the group mean stimulated MNI coordinate across participants, showing that TMS was applied to the right ANG/IPS, a region normally activated during bilateral attention ($x = 36$, $y = -64$, $z = 42$). The circle was positioned by centering a 1 cm radius sphere one radius distance inward from the mean scalp coordinates along a trajectory perpendicular to the angle of orientation of the TMS coil.

Area	MNI coordinates			Z-value
	x	y	z	
Visual cortex				
R	44	-64	-12	4.00
L	-40	-60	-12	3.90
Angular gyrus/intra-parietal sulcus				
R	34	-64	40	3.07
L	-44	-62	50	3.02
Superior parietal lobe				
R	32	-52	62	3.20
L	-28	-56	50	2.95
Precentral gyrus				
L	-32	0	36	2.94

Table B.1: Network of brain regions significantly activated by the extinction task. Coordinates are the peak voxel for the contrast [Left/Bilateral/Right – Catch] averaged across the baseline and Post-TMS scan sessions.

B.3.2.2 Effect of TMS on the inter-hemispheric balance

Whole-brain voxel-wise analysis ($Z \geq 2.3$, corrected cluster significance threshold of $p = 0.05$) revealed no overall increase or decrease in activity anywhere in the brain between the baseline and the post-TMS scan sessions. However, the goal of the study was to test the more specific hypothesis that right ANG/IPS TMS would induce a leftward shift in the relative balance of activity between the left and right parietal cortices. The statistical map of the extinction task network computed for the baseline scan suggested there was relatively stronger activity in the right parietal cortex, whereas the same map computed for the post-TMS scan suggested the opposite (Figure B.4).

In order to formally assess this, we computed the relative difference in activity between homologous voxels of the two hemispheres for the baseline and the post-TMS scans, and tested for a significant change in functional lateralisation caused by TMS. Voxelwise functional lateralisation maps were obtained for each scan session using a

similar approach to other fMRI studies of functional lateralisation (Macaluso & Patria 2007; Shulman *et al.* 2010; Agcaoglu *et al.* 2015). To accurately map homologous voxels across hemispheres, we used a simple left-right flip of the images, followed by a non-linear registration of the (flipped) left hemisphere to the (non-flipped) right hemisphere (Jenkinson *et al.* 2002). We then generated functional lateralisation maps by comparing the non-flipped left hemisphere to the flipped right hemisphere images. For each voxel of the right hemisphere, the lateralisation index was calculated as the difference in effect size relative to its contralateral homologous voxel (ie: right-left hemisphere) for the contrast [Left/Bilateral/Right - Catch trials]. At the group level, baseline and post-TMS functional lateralisation maps were masked to include only grey matter and were compared using paired sample t-tests across participants for each voxel. The resulting statistical maps were thresholded using clusters determined by $t(11) \geq 2.2$ (corresponding to $p \leq .025$) and a corrected cluster significance threshold of $p = 0.05$.

Right ANG/IPS TMS caused a significant leftward shift in hemispheric balance (functional lateralisation) that was confined to the stimulated angular gyrus and adjacent intraparietal sulcus (peak voxel: $X = 54$, $Y = -54$, $Z = 40$, $t(11) = 6.14$, $p < 10^{-4}$; [Figure B.4B](#), C). No brain region showed a rightward shift in hemispheric balance after TMS. For illustration purposes, [Figure B.5](#) shows the effect of TMS on brain activity for each trial type for the region that showed the significant leftward shift in inter-hemispheric dominance.

B.3.2.3 Inter-individual variation in the magnitude of the TMS effect

To visualise inter-individual variation in response to TMS, we extracted the inter-hemispheric balance parameter estimates for the baseline and post-TMS scanning sessions from the regions of significant TMS effect, and asked whether the strength




Figure B.4: Effect of rTMS on inter-hemispheric balance during bilateral attention. **A.** Network activated during bilateral attention before and after TMS. Figure shows regions activated during the extinction task, separately for the baseline (blue) and post-TMS (red) scan sessions. Colour bar represents Z statistics at each voxel. Voxelwise statistics represent the contrast [Left/Bilateral/Right – Catch trials] with a voxel height threshold of $Z > 1.96$ and a corrected cluster extent significance threshold of $p < 0.05$. At this threshold, the relative difference in the strength of activity across hemispheres is visible. At baseline there is more supra-threshold activity in the right versus left parietal cortex. Post-TMS this is inverted. **B.** Right parietal TMS shifts the inter-hemispheric balance of activity leftward. Voxelwise analysis of functional lateralisation within the extinction task network confirmed the apparent difference visible in A) using a direct and unbiased statistical comparison across the two hemispheres. Homologous voxels were co-registered onto the right hemisphere and compared statistically. Map shows the only region in which there was a significant leftward shift in the relative balance of task activity (right-left hemisphere) after TMS: ANG and the adjacent IPS (stimulated cortex). The white circle illustrates the group mean stimulated MNI coordinate across participants ($x = 36, y = -64, z = 42$). The circle was positioned by centering a 1cm radius sphere one radius distance inward from the mean scalp coordinates along a trajectory perpendicular to the angle of orientation of the TMS coil. The slice $y = -54$ was added for better visualisation of the TMS effect. Colour bar represents t statistics at each voxel. **C.** Extracted parameter estimates from the region of significant TMS effect. Data extracted from the region of significant TMS effect (as in B) illustrate the inter-hemispheric balance parameter (in arbitrary units, a.u.) from the baseline and post-TMS sessions. R = right hemisphere, L = left hemisphere




Figure B.5: Effect of TMS on brain activity in the region showing a significant leftward shift in inter-hemispheric balance.

of right parietal dominance at baseline predicted the magnitude of change in balance caused by TMS.

In psychology, this relation is known as *the law of initial value* (Jin 1992) and is associated with two methodological concerns known as *mathematical coupling* and *regression towards the mean* (Blance *et al.* 2005; Tu & Gilthorpe 2007). Regression towards the mean occurs when a variable is measured repeatedly over time and refers to the fact that if, because of measurement noise and/or physiological variation, a first measurement is abnormally far from its mean value, then the second measurement is more likely to be closer to this mean value. Because of this statistical phenomenon participants with extreme baseline values are more likely to exhibit a greater change than the one with baseline values closer to the mean value. However, the counterbalancing of the two scanning sessions ruled out any systematic bias introduced by regression towards the mean because its contribution was randomised across the two orders (Blance *et al.* 2005). Mathematical coupling occurs in a correlation when one variable — in this case the change in inter-hemispheric balance parameter — contains the whole or part of the other — in this case the baseline inter-hemispheric balance parameter. As a result, the null hypothesis that the correlation coefficient is equal to zero is no longer appropriate and some adjustment is needed. When testing the relation between baseline and change in parietal inter-hemispheric balance




Figure B.6: Baseline-dependency of the effect of TMS on parietal inter-hemispheric balance. The graph illustrates the variability across individuals in response to TMS. Data are extracted from the region of significant TMS effect (as in [Figure B.4](#) and [B.5](#)). The x-axis represents the average of baseline and post-tDCS scans in accordance to Oldham’s method (Oldham 1962).

parameter, we found a Pearson correlation coefficient of $r_{(12)} = -0.92$ suggesting a very strong negative correlation between the two. However for the reason highlighted above this value is artificially inflated because of mathematical coupling. In order to account for this phenomenon, we applied Oldham’s method (Oldham 1962; Tu & Gilthorpe 2007) and tested the correlation of the average inter-hemispheric balance parameter across the baseline and post-TMS scanning sessions versus the change in inter-hemispheric balance parameter. When doing so, we found a marginally significant correlation ($r_{(12)} = -0.56, p = 0.056$) suggesting that, across participants, the stronger the right hemisphere dominance at baseline, the greater the leftward shift in hemispheric balance after TMS ([Figure B.6](#)).

B.4 Discussion

This study aimed to test a hypothesis generated from theoretical models of the pathophysiology of visual extinction and neglect: that bilateral attentional function in the healthy brain depends causally on the balance of functional activity between left and right parietal cortex (Kinsbourne 1976; Cohen *et al.* 1994). To test this, participants performed a target detection task in which they were required to maintain central fixation while covertly monitoring both hemifields simultaneously to detect and report the location of an upcoming target stimulus. Low-frequency TMS was applied to transiently inhibit the right ANG/IPS. Experiment 1 assessed the behavioural effects of this stimulation protocol under psychophysically controlled conditions. Analysis of errors on bilateral trials revealed that, prior to TMS, participants exhibited an attentional bias towards the left hemifield (pseudoneglect)(Figure B.2). TMS shifted the spatial distribution of errors from left to right, simulating visual extinction and abolishing pseudoneglect. Experiment 2 then assessed the effect of TMS on functional brain activity during bilateral attention. TMS did not cause an overall increase or decrease in activity anywhere in the brain. Rather, TMS shifted the *relative balance* of inter-hemispheric activity from right to left, an effect that was confined to parietal cortex. In combination, across the two experiments, the coherent direction of change in the TMS-induced behavioural (left-to-right) and functional (right-to-left) effects supports the hypothesis that bilateral spatial attention causally depends on parietal inter-hemispheric balance.

Given the inability to obtain precise psychophysical control over visual stimuli within the scanner, a behavioural experiment was first conducted outside the scanner in a subset of participants to confirm that stimulation of the right ANG/IPS would induce an extinction-like behavioural effect (for similar logic in an online TMS/fMRI study see Ruff *et al.*, 2006). The same visual target detection task was used as in

the subsequent fMRI experiment, except that the task in the laboratory was better sensitised to detect a behavioural effect of TMS. Trial pacing was faster (mean ITI of 3 versus 7 seconds), requiring participants to perform a larger number of trials within the same 15 minute time period (144 versus 100), and there was no delay between the offset of TMS and the onset of task performance. Critically, visual luminance could be precisely controlled, ensuring that the staircase procedure more accurately estimated participant's psychophysical detection threshold. Under these conditions, TMS induced a subtle but measurable behavioural effect. Analysis of the spatial distribution of errors on bilateral trials showed that at baseline, prior to stimulation, participants exhibited an attentional bias towards the left hemifield, a phenomenon known as pseudoneglect. Pseudoneglect is a leftward bias in spatial attention reliably observed in healthy individuals across a range of tasks (Bowers & Heilman 1980; Nicholls *et al.* 1999; Jewell & McCourt 2000; Mattingley *et al.* 2004). Taken together with the higher prevalence of spatial neglect after right than left hemisphere lesions, the phenomenon of pseudoneglect in healthy individuals is thought to reflect right hemisphere dominance (or even lateralisation) of some visual attention processes in the general population (çiçek *et al.* 2009; Benwell *et al.* 2014; Le *et al.* 2015). After stimulation, pseudoneglect was abolished with participants now tending to report a target on the right and omitting the target on left (Figure B.2), thus simulating clinical visual extinction. Thus the behavioural TMS experiment established that right ANG/IPS has a causal role in bilateral attention, replicating previous work (Hilgetag *et al.* 2001).

Previous behavioural TMS studies had reported that stimulation of right inferior parietal cortex could induce extinction-like behaviour, which was presumed to reflect TMS-induced changes in the balance of functional activity between left and right parietal cortex (Pascual-Leone *et al.* 1994; Hilgetag *et al.* 2001; Dambeck *et al.*

2006; Meister *et al.* 2006). The present study was designed to test this hypothesis directly, by imaging the causal impact of stimulation on bilateral attention networks using whole-brain fMRI. The extinction task activated an expected bilateral network of parietal, frontal and visual cortical areas (Figure B.3, Table B.1), including the region of right ANG/IPS that was the target for TMS. Whole-brain voxelwise analysis showed that stimulation did not cause an overall increase or decrease in activity. Rather, stimulation shifted the *relative balance* of activity between the right and left hemisphere, inducing a relative leftward shift selectively in the ANG/IPS while participants were covertly monitoring both hemifields (Figure B.4). An advantage of this analysis approach is that it tests directly the central claim of the original inter-hemispheric rivalry model - that it is the *relative* balance between (and not the overall level of activity within) left and right parietal cortex that determines the spatial allocation of attention (Kinsbourne 1976). The counterbalanced order of baseline and post-TMS scan sessions rules out alternative artefactual explanations of this effect, such as non-specific changes in arousal after TMS, task learning across scan sessions, or fatigue from increasing time-on-task.

A key feature of the target detection task was that, on every trial, participants were required to monitor both hemifields simultaneously in anticipation of an upcoming target. Thus, every trial started with a cue period (2-4 seconds) that required bilateral spatial attention. Whilst the rapid event-related fMRI design did not allow for cue- and target-related periods to be dissociated from one another within a trial, two aspects of the results suggest that stimulation most likely affected this bilaterally-directed cue-period activity, rather than subsequent target-related processing. First, stimulation changed the pattern of brain activity similarly across all trial types (right, left, bilateral), consistent with a similar function being disrupted on all trials (Figure B.5). Second, examination of the timecourse of fMRI signal change within the




Figure B.7: Timecourse of TMS-induced change in inter-hemispheric balance. Graphs show the average timecourse of activation change within a trial for the baseline (A) and post-TMS (B) sessions extracted from the region of significant TMS effect (Figure B.4B). Timecourses are aligned on cue onset, and show a change in functional lateralisation. When the same data are aligned instead to target rather than cue onset there is no such difference. This suggests that stimulation disrupted bilateral attention distribution during the cue period of each trial.

region of TMS effect (Figure B.7) clearly shows a change in inter-hemispheric balance when the data are aligned at cue onset; no signal difference is observed when the same data are aligned on target onset (data not shown). Hence, the fMRI data suggest that right parietal interference (with TMS, but possibly also in the case of a lesion) likely induces a rightward bias in top-down preparatory spatial attention signals, rather than altering target-related processing, as might alternatively account for extinction. The behavioural cost of such an induced rightward functional bias is likely to be most readily apparent on bilateral target trials, as these are the most attentionally demanding and therefore the most behaviourally sensitive.

By contrast with the leftward spatial bias observed at baseline in Experiment 1 (and subsequently abolished by TMS), no such pseudoneglect was present in the fMRI study. The visual environment in the MRI scanner precluded precise psychophysical control over luminance parameters, and hence the perceptibility of stimuli. Hence, it was not possible to titrate stimuli to psychophysical threshold inside the scanner with the same precision as outside the scanner, meaning that the fMRI experiment

was necessarily less sensitive to subtle behavioural change. This was already clear at the pilot stage. The fact that the same mean stimulus size was used in both experiments (ie: 12 pixels), but performance was on average 9% higher at baseline in the fMRI study (compare [Figure B.2](#) and [B.4](#)) testifies to this. We believe that the absence of baseline pseudoneglect inside the scanner, despite its presence outside the scanner in a subset of the same participants, reflects this reduced measurement sensitivity. In addition, the rapid event-related design required a relatively long inter-trial interval (6-8 sec). This, combined with the need to complete scanning within 15 minutes (to maximise the likely window of decaying TMS effect), entailed fewer trials in the scanner. This reduced power compared to the equivalent behavioural experiment outside the scanner (reduction from 144 to 100 trials). More importantly, slowing the pace of task performance made the task less demanding, further reducing behavioural sensitivity. Hence, in summary, the fMRI experiment was optimised to detect the physiological impact of TMS on the bilateral attention network, but not its behavioural consequences.

The absence of pseudoneglect in the scanner environment meant there was no baseline behavioural phenomenon for TMS to change. Hence, unlike in Experiment 1, TMS in the scanner had no measurable effect on extinction task performance. One benefit of this is that it enabled us to attribute any observed changes in functional brain activity to a direct causal effect of TMS, without any contamination by behavioural change. The absence of a behavioural effect in the scanner does, however, raise the question of the functional significance of the observed TMS-induced change in inter-hemispheric balance. There are at least three possible interpretations: 1) the changes are epiphenomenal; 2) they reflect a compensatory response to stimulation; or 3) they reflect neural changes likely to cause extinction-like behaviour if measured under psychophysically controlled conditions. We think the observed changes

are unlikely to be epiphenomenal. The fMRI effect cannot be explained as a simple passive consequence of the removal of transcallosal inhibition, whereby inhibitory stimulation to right ANG/IPS would cause an automatic increase in excitation in left parietal cortex. Such an effect, if present, would be common to both the implicit resting baseline and task-related periods, and since fMRI contrasts subtract across these conditions, this would cancel out and hence not be detectable. Consequently, the change in inter-hemispheric balance indicates that TMS modulated the *functioning* of right ANG/IPS, rather than simply changing cortical excitability and leaving function unaffected. The second possibility is that the changes reflect compensation in response to TMS-induced functional interference. We cannot rule this out. Note the necessary delay in the TMS/fMRI experiment between the end of stimulation and the onset of MRI data acquisition (about 4 minutes). Any behavioural interference effects of offline TMS are likely to be strongest in this immediate post-stimulation period, before the system has had much time to recover. In this period behavior could not be measured. The combination of this 4-minute delay, together with an offline stimulation protocol, may have provided sufficient opportunity for the brain to compensate sufficiently for the stimulation so as to maintain task performance despite neural activation change (see for example O'Shea *et al.* 2007). However, given that Experiment 1 showed no effect of TMS on any task measure other than pseudoneglect, and given that pseudoneglect was already absent during the baseline scan, we do not think a 'compensatory' account makes explanatory sense. Our preferred, parsimonious interpretation is that the coherent direction of TMS-induced behavioural change (rightward shift) in Experiment 1, and functional change (leftward shift) in Experiment 2, is consistent with inter-hemispheric competition models of visual extinction (Kinsbourne 1976; Szczepanski *et al.* 2010; Szczepanski & Kastner 2013). Hence, we speculate that the TMS-induced changes in brain activity observed in Experiment 2 are likely to causally underwrite the TMS-induced behavioural effect

observed in Experiment 1, but the latter requires psychophysically controlled measurement conditions to be detectable. In brief, we speculate that these TMS-induced neural changes may be causally necessary to induce an extinction-like behavioural change, but they are not sufficient.

B.5 Conclusions

The present study is the first, to our knowledge, to directly test the inter-hemispheric balance model of visual extinction in the healthy brain, using a task similar in kind to that used in the clinic. Whilst the behavioural and functional effects of TMS were observed in two separate experiments, the direction of TMS-induced change was consistent across the two studies (i.e. leftward shift in brain activity, rightward shift in behaviour), and occurred in individuals who had participated in both experiments. Hence, we speculate that the TMS-induced change in inter-hemispheric balance observed here is likely, under behaviourally demanding and psychophysically controlled conditions, to induce extinction-like behavioural interference, similar to that observed outside the scanner. These results support the claim of theoretical models that bilateral attention in the healthy brain depends critically on the balance of functional brain activity between left and right parietal cortex.

Appendix C

Supplementary information from Chapter 6

This appendix reports additional analyses run to evaluate the robustness of the MRS results to different analysis methods.

Older adults have lower GABAergic inhibition (section 6.3.4)

Expressing GABA and Glx levels as ratios to tCr rather than absolute concentrations (as reported in the main text of this thesis) did not change the pattern of age-related decrease in V1 GABA:

- Age $\times$ M1 GABA:tCr: $r_{(22)} = -0.46$, $p = 0.03$, 1-tail, controlling for [GM]_{M1}, [WM]_{M1} and M1 Glx:tCr
- Age $\times$ V1 GABA:tCr: $r_{(16)} = -0.49$, $p = 0.04$, 1-tail, controlling for [GM]_{V1}, [WM]_{V1} and V1 Glx:tCr
- Age $\times$ M1 Glx:tCr: $r_{(22)} = -0.04$, $p = 0.86$, controlling for [GM]_{M1}, [WM]_{M1}

and M1 GABA:tCr

- Age $\times$ V1 Glx:tCr: $r_{(16)} = -0.09$, $p = 0.76$, controlling for [GM]_{V1}, [WM]_{V1} and V1 GABA:tCr

Lower M1 GABA is associated with more persistent after-effect (section 6.3.5)

Expressing GABA and Glx levels as ratios to tCr rather than absolute concentrations (as reported in the main text of this thesis) did not change the relationship between M1 GABA and AE 24 hours after the adaptation:

- M1 GABA:tCr $\times$ AE_{24hrs}: $r_{(22)} = 0.79$, $p < 10^{-4}$, 1-tail, controlling for [GM]_{M1}, [WM]_{M1} and M1 Glx:tCr
- M1 Glx:tCr $\times$ AE_{24hrs}: $r_{(22)} = 0.14$, $p = 0.56$, controlling for [GM]_{M1}, [WM]_{M1} and M1 GABA:tCr
- V1 GABA:tCr $\times$ AE_{24hrs}: $r_{(16)} = 0.08$, $p = 0.79$, controlling for [GM]_{V1}, [WM]_{V1} and V1 Glx:tCr
- V1 Glx:tCr $\times$ AE_{24hrs}: $r_{(16)} = 0.43$, $p = 0.13$, controlling for [GM]_{V1}, [WM]_{V1} and V1 GABA:tCr
- Neurochemical specificity: $Z = 2.84$, $p = 0.005$
- Anatomical specificity: $Z = 2.72$, $p = 0.006$

Mediation analysis (section 6.3.6)

Expressing GABA and Glx levels as ratios to tCr rather than absolute concentrations (as reported in the main text of this thesis) did not change the results of the mediation analysis:

- $c = -0.76$, $p < 10^{-3}$
- $c' = -0.33$, $p = 0.08$
- $a_1 = -0.66$, $p = 0.001$
- $b_1 = 0.66$, $p = 0.003$
- $ab_1 = 0.44$, $p < 10^{-3}$
- $a_2 = 0.02$, $p = 0.47$
- $b_2 = 0.11$, $p = 0.33$
- $ab_2 = 0.00$, $p = 0.40$

Bibliography

Adamson, J., Beswick, A. & Ebrahim, S. Is stroke the most common cause of disability? *Journal of Stroke and Cerebrovascular Diseases* **13**, 171–177 (2004).

Agcaoglu, O., Miller, R., Mayer, A. R., Hugdahl, K. & Calhoun, V. D. Lateralization of resting state networks and relationship to age and gender. *Neuroimage* **104**, 310–325 (2015).

Aimola, L., Rogers, G., Kerkhoff, G., Smith, D. T. & Schenk, T. Visuomotor adaptation is impaired in patients with unilateral neglect. *Neuropsychologia* **50**, 1158–1163 (2012).

Akshoomoff, N. A. & Courchesne, E. ERP evidence for a shifting attention deficit in patients with damage to the cerebellum. *Journal of Cognitive Neuroscience* **6**, 388–399 (1994).

Alais, D. & Burr, D. The ventriloquist effect results from near-optimal bimodal integration. *Current biology* **14**, 257–262 (2004).

Albert, S. T. & Shadmehr, R. Estimating properties of the fast and slow adaptive processes during sensorimotor adaptation. *Journal of neurophysiology*, jn-00197 (2017).

Allen, G. & Courchesne, E. The cerebellum and non-motor function: clinical implications. *Molecular psychiatry* **3**, 207–210 (1998).

Allen, G., Buxton, R. B., Wong, E. C. & Courchesne, E. Attentional activation of the cerebellum independent of motor involvement. *Science* **275**, 1940–1943 (1997).

Amadi, U., Allman, C., Johansen-Berg, H. & Stagg, C. J. The homeostatic interaction between anodal transcranial direct current stimulation and motor learning in humans is related to GABA A activity. *Brain stimulation* **8**, 898–905 (2015).

Amaro, E. & Barker, G. J. Study design in fMRI: basic principles. *Brain and cognition* **60**, 220–232 (2006).

Ammann, C., Spampinato, D. & Márquez-Ruiz, J. Modulating motor learning through transcranial direct-current stimulation: an integrative view. *Frontiers in psychology* **7** (2016).

Anderson, M. J. & Robinson, J. Permutation tests for linear models. *Australian & New Zealand Journal of Statistics* **43**, 75–88 (2001).

- Andersson, J. L., Skare, S. & Ashburner, J. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *Neuroimage* **20**, 870–888 (2003).
- Andersson, J. L. & Sotiropoulos, S. N. An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. *Neuroimage* **125**, 1063–1078 (2016).
- Andersson, J. L., Jenkinson, M., Smith, S., *et al.* Non-linear registration, aka Spatial normalisation FMRIB technical report TR07JA2. *FMRIB Analysis Group of the University of Oxford* **2** (2007).
- Anguera, J. A., Reuter-Lorenz, P. A., Willingham, D. T. & Seidler, R. D. Failure to engage spatial working memory contributes to age-related declines in visuomotor learning. *Journal of cognitive neuroscience* **23**, 11–25 (2011).
- Antal, A., Nitsche, M. A., Kincses, T. Z., Lampe, C. & Paulus, W. No correlation between oving phosphene and motor thresholds: a transcranial magnetic stimulation study. *Neuroreport* **15**, 297–302 (2004).
- Antal, A., Terney, D., Kühnl, S. & Paulus, W. Anodal transcranial direct current stimulation of the motor cortex ameliorates chronic pain and reduces short intracortical inhibition. *Journal of pain and symptom management* **39**, 890–903 (2010).
- Antonenko, D., Schubert, F., Bohm, F., Ittermann, B., Aydin, S., Hayek, D., Grittnet, U. & Flöel, A. tDCS-Induced Modulation of GABA Levels and Resting-State Functional Connectivity in Older Adults. *Journal of Neuroscience* **37**, 4065–4073 (2017).
- Appelros, P., Karlsson, G. M., Seiger, A., Nydevik, I., *et al.* Neglect and anosognosia after first-ever stroke: incidence and relationship to disability. *Journal of Rehabilitation Medicine* **34**, 215–220 (2002).
- Ashbridge, E., Walsh, V. & Cowey, A. Temporal aspects of visual search studied by transcranial magnetic stimulation. *Neuropsychologia* **35**, 1121–1131 (1997).
- Assaf, Y., Johansen-Berg, H. & de Schotten, M. T. The Role Of Diffusion MRI In Neuroscience. *bioRxiv*, 140459 (2017).
- Azouvi, P. Functional consequences and awareness of unilateral neglect: study of an evaluation scale. *Neuropsychological Rehabilitation* **6**, 133–150 (1996).
- Azouvi, P. The ecological assessment of unilateral neglect. *Annals of physical and rehabilitation medicine* **60**, 186–190 (2017).
- Azouvi, P., Jacquin-Courtois, S. & Luauté, J. Rehabilitation of unilateral neglect: Evidence-based medicine. *Annals of physical and rehabilitation medicine* (2016).
- Azouvi, P., Olivier, S., De Montety, G., Samuel, C., Louis-Dreyfus, A. & Tesio, L. Behavioral assessment of unilateral neglect: study of the psychometric properties of the Catherine Bergego Scale. *Archives of physical medicine and rehabilitation* **84**, 51–57 (2003).

- Azouvi, P., Samuel, C., Louis-Dreyfus, A., Bernati, T., Bartolomeo, P., Beis, J., Chokron, S., Leclercq, M., Marchal, F., Martin, Y., *et al.* Sensitivity of clinical and behavioural tests of spatial neglect after right hemisphere stroke. *Journal of Neurology, Neurosurgery & Psychiatry* **73**, 160–166 (2002).
- Bachtiar, V. & Stagg, C. The role of inhibition in human motor cortical plasticity. *Neuroscience* **278**, 93–104 (2014).
- Bachtiar, V., Near, J., Johansen-Berg, H. & Stagg, C. J. Modulation of GABA and resting state functional connectivity by transcranial direct current stimulation. *Elife* **4**, e08789 (2015).
- Baizer, J. S., Kralj-Hans, I. & Glickstein, M. Cerebellar lesions and prism adaptation in macaque monkeys. *Journal of neurophysiology* **81**, 1960–1965 (1999).
- Balkowiec, A. & Katz, D. M. Cellular mechanisms regulating activity-dependent release of native brain-derived neurotrophic factor from hippocampal neurons. *Journal of Neuroscience* **22**, 10399–10407 (2002).
- Bank, B., Emir, U., Boer, V., Asten, J., Maas, M., Wijnen, J., Kan, H., Oz, G., Klomp, D. & Scheenen, T. Multi-center reproducibility of neurochemical profiles in the human brain at 7 T. *NMR in Biomedicine* **28**, 306–316 (2015).
- Barer, D. The influence of visual and tactile inattention on predictions for recovery from acute stroke. *QJM: An International Journal of Medicine* **74**, 21–32 (1990).
- Barker, A. T., Jalinous, R. & Freeston, I. L. Non-invasive magnetic stimulation of human motor cortex. *The Lancet* **325**, 1106–1107 (1985).
- Barrett, A., Goedert, K. M. & Basso, J. C. Prism adaptation for spatial neglect after stroke: translational practice gaps. *Nature Reviews Neurology* **8**, 567–577 (2012).
- Barrett, A. M., Buxbaum, L. J., Coslett, H. B., Edwards, E., Heilman, K. M., Hillis, A. E., Milberg, W. P. & Robertson, I. H. Cognitive rehabilitation interventions for neglect and related disorders: moving from bench to bedside in stroke patients. *Journal of cognitive neuroscience* **18**, 1223–1236 (2006).
- Barron, H., Vogels, T., Emir, U., Makin, T., O’Shea, J., Clare, S., Jbabdi, S., Dolan, R. & Behrens, T. Unmasking latent inhibitory connections in human cortex to reveal dormant cortical memories. *Neuron* **90**, 191–203 (2016).
- Bays, P. M., Singh-Curry, V., Gorgoraptis, N., Driver, J. & Husain, M. Integration of goal-and stimulus-related visual signals revealed by damage to human parietal cortex. *Journal of Neuroscience* **30**, 5968–5978 (2010).
- Beaulieu, C. The biological basis of diffusion anisotropy. *Diffusion MRI: From quantitative measurement to in vivo neuroanatomy*, 105–126 (2009).
- Becker, E. & Karnath, H.-O. Incidence of visual extinction after left versus right hemisphere stroke. *Stroke* **38**, 3172–3174 (2007).
- Beckmann, C. F. Modelling with independent components. *Neuroimage* **62**, 891–901 (2012).

- Beckmann, C. F., DeLuca, M., Devlin, J. T. & Smith, S. M. Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* **360**, 1001–1013 (2005).
- Beckmann, C. F., Mackay, C. E., Filippini, N. & Smith, S. M. Group comparison of resting-state fMRI data using multi-subject ICA and dual regression. *Neuroimage* **47**, S148 (2009).
- Beckmann, C. F. & Smith, S. M. Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE transactions on medical imaging* **23**, 137–152 (2004).
- Beckmann, M., Johansen-Berg, H. & Rushworth, M. F. Connectivity-based parcellation of human cingulate cortex and its relation to functional specialization. *Journal of Neuroscience* **29**, 1175–1190 (2009).
- Bedford, F. L. Constraints on learning new mappings between perceptual dimensions. *Journal of Experimental Psychology: Human Perception and Performance* **15**, 232 (1989).
- Bedford, F. L. Perceptual and cognitive spatial learning. *Journal of Experimental Psychology: Human Perception and Performance* **19**, 517 (1993).
- Behrens, T. E., Berg, H. J., Jbabdi, S., Rushworth, M. F. & Woolrich, M. W. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *Neuroimage* **34**, 144–155 (2007).
- Behrens, T. E., Woolrich, M. W., Jenkinson, M., Johansen-Berg, H., Nunes, R. G., Clare, S., Matthews, P. M., Brady, J. M. & Smith, S. M. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic resonance in medicine* **50**, 1077–1088 (2003).
- Behrens, T. E., Woolrich, M. W., Walton, M. E. & Rushworth, M. F. Learning the value of information in an uncertain world. *Nature neuroscience* **10**, 1214–1221 (2007).
- Benwell, C. S., Harvey, M. & Thut, G. On the neural origin of pseudoneglect: EEG-correlates of shifts in line bisection performance with manipulation of line length. *Neuroimage* **86**, 370–380 (2014).
- Bergego, C., Azouvi, P., Samuel, C., Marchal, F., Louis-Dreyfus, A., Jokic, C., Morin, L., Renard, C., Pradat-Diehl, P. & Deloche, G. *Validation d'une échelle d'évaluation fonctionnelle de l'héminégligence dans la vie quotidienne: l'échelle CB Annales de réadaptation et de médecine physique* **38** (1995), 183–189.
- Bernacchia, A., Seo, H., Lee, D. & Wang, X.-J. A reservoir of time constants for memory traces in cortical neurons. *Nature neuroscience* **14**, 366–372 (2011).
- Berniker, M. & Kording, K. Estimating the sources of motor errors for adaptation and generalization. *Nature neuroscience* **11**, 1454–1461 (2008).
- Berniker, M. & Kording, K. Bayesian approaches to sensory integration for motor control. *Wiley Interdisciplinary Reviews: Cognitive Science* **2**, 419–428 (2011).

- Bestmann, S. & Walsh, V. Transcranial electrical stimulation. *Current Biology* **27**, R1258–R1262 (2017).
- Bickerton, W. L., Samson, D., Williamson, J. & Humphreys, G. W. Separating forms of neglect using the Apples Test: validation and functional prediction in chronic and acute stroke. *Neuropsychology* **25**, 567 (2011).
- Bikson, M., Grossman, P., Thomas, C., Zannou, A. L., Jiang, J., Adnan, T., Mourdoukoutas, A. P., Kronberg, G., Truong, D., Boggio, P., *et al.* Safety of transcranial direct current stimulation: evidence based update 2016. *Brain stimulation* **9**, 641–661 (2016).
- Binder, J., Marshall, R., Lazar, R., Benjamin, J. & Mohr, J. Distinct syndromes of hemineglect. *Archives of Neurology* **49**, 1187–1194 (1992).
- Bindman, L. J., Lippold, O. & Redfearn, J. Long-lasting changes in the level of the electrical activity of the cerebral cortex produced by polarizing currents. *Nature* **196**, 584–585 (1962).
- Bindman, L. J., Lippold, O. & Redfearn, J. The action of brief polarizing currents on the cerebral cortex of the rat (1) during current flow and (2) in the production of long-lasting after-effects. *The Journal of physiology* **172**, 369–382 (1964).
- Birn, R. M., Diamond, J. B., Smith, M. A. & Bandettini, P. A. Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. *Neuroimage* **31**, 1536–1548 (2006).
- Blance, A., Tu, Y.-K. & Gilthorpe, M. S. A multilevel modelling solution to mathematical coupling. *Statistical methods in medical research* **14**, 553–565 (2005).
- Bliss, T. V. & Collingridge, G. L. A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* **361**, 31 (1993).
- Block, H. & Celnik, P. Stimulating the cerebellum affects visuomotor adaptation but not intermanual transfer of learning. *The Cerebellum* **12**, 781–793 (2013).
- Bock, O. Components of sensorimotor adaptation in young and elderly subjects. *Experimental Brain Research* **160**, 259–263 (2005).
- Bonato, M. Neglect and extinction depend greatly on task demands: a review. *Frontiers in Human Neuroscience* **6** (2012).
- Bonato, M., Priftis, K., Umiltà, C. & Zorzi, M. Computer-based attention-demanding testing unveils severe neglect in apparently intact patients. *Behavioural Neurology* **26**, 179–181 (2013).
- Boroojerdi, B., Prager, A., Muellbacher, W. & Cohen, L. G. Reduction of human visual cortex excitability using 1-Hz transcranial magnetic stimulation. *Neurology* **54**, 1529–1531 (2000).
- Boroojerdi, B., Meister, I. G., Foltys, H., Sparing, R., Cohen, L. G. & Töpper, R. Visual and motor cortex excitability: a transcranial magnetic stimulation study. *Clinical Neurophysiology* **113**, 1501–1504 (2002).

- Bowen, A., Hazelton, C., Pollock, A. & Lincoln, N. B. Cognitive rehabilitation for spatial neglect following stroke. *The Cochrane Library* (2013).
- Bowen, A., Lincoln, N. & Dewey, M. Cognitive rehabilitation for spatial neglect following stroke. *Cochrane Database Syst Rev* **2** (2007).
- Bowers, D. & Heilman, K. M. Pseudoneglect: effects of hemispace on a tactile line bisection task. *Neuropsychologia* **18**, 491–498 (1980).
- Bozzali, M., Mastropasqua, C., Cercignani, M., Giulietti, G., Bonni, S., Caltagirone, C. & Koch, G. Microstructural damage of the posterior corpus callosum contributes to the clinical severity of neglect. *PloS one* **7**, e48079 (2012).
- Bramham, C. R. & Messaoudi, E. BDNF function in adult synaptic plasticity: the synaptic consolidation hypothesis. *Progress in neurobiology* **76**, 99–125 (2005).
- Bressler, S. L., Tang, W., Sylvester, C. M., Shulman, G. L. & Corbetta, M. Top-down control of human visual cortex by frontal and parietal cortex in anticipatory visual spatial attention. *Journal of Neuroscience* **28**, 10056–10061 (2008).
- Broucke, M. E. A State Space Model of a Visuomotor Rotation Experiment. *bioRxiv*, 197483 (2017).
- Brown, R. D., Whisnant, J. P., Sicks, J., O’Fallon, W. M. & Wiebers, D. O. Stroke incidence, prevalence, and survival: secular trends in Rochester, Minnesota, through 1989. *Stroke* **27**, 373–380 (1996).
- Buch, E. R., Young, S. & Contreras-Vidal, J. L. Visuomotor adaptation in normal aging. *Learning & memory* **10**, 55–63 (2003).
- Burge, J., Girshick, A. R. & Banks, M. S. Visual–haptic adaptation is determined by relative reliability. *Journal of Neuroscience* **30**, 7714–7721 (2010).
- Bütefisch, C. M., Davis, B. C., Wise, S. P., Sawaki, L., Kopylev, L., Classen, J. & Cohen, L. G. Mechanisms of use-dependent plasticity in the human motor cortex. *Proceedings of the national academy of sciences* **97**, 3661–3665 (2000).
- Buxton, R. B., Uludag, K., Dubowitz, D. J. & Liu, T. T. Modeling the hemodynamic response to brain activation. *Neuroimage* **23**, S220–S233 (2004).
- C. Justus, T. & Ivry, R. B. The cognitive neuropsychology of the cerebellum. *International Review of Psychiatry* **13**, 276–282 (2001).
- Calzolari, E., Bolognini, N., Casati, C., Marzoli, S. B. & Vallar, G. Restoring abnormal aftereffects of prismatic adaptation through neuromodulation. *Neuropsychologia* **74**, 162–169 (2015).
- Canavan, A., Passingham, R., Marsden, C., Quinn, N., Wyke, M. & Polkey, C. Prism adaptation and other tasks involving spatial abilities in patients with Parkinson’s disease, patients with frontal lobe lesions and patients with unilateral temporal lobectomies. *Neuropsychologia* **28**, 969–984 (1990).
- Carroll, T. J., Poh, E. & de Rugy, A. New visuomotor maps are immediately available to the opposite limb. *Journal of neurophysiology* **111**, 2232–2243 (2014).

- Carter, A. R., Astafiev, S. V., Lang, C. E., Connor, L. T., Rengachary, J., Strube, M. J., Pope, D. L., Shulman, G. L. & Corbetta, M. Resting interhemispheric functional magnetic resonance imaging connectivity predicts performance after stroke. *Annals of neurology* **67**, 365–375 (2010).
- Carter, A. R., McAvoy, M. P., Siegel, J. S., Hong, X., Astafiev, S. V., Rengachary, J., Zinn, K., Metcalf, N. V., Shulman, G. L. & Corbetta, M. Differential white matter involvement associated with distinct visuospatial deficits after right hemisphere stroke. *Cortex* **88**, 81–97 (2017).
- Catani, M. & Thiebaut de Schotten, M. *Atlas of human brain connections* (Oxford University Press, 2012).
- Cavassila, S., Deval, S., Huegen, C., Van Ormondt, D. & Graveron-Demilly, D. Cramér–Rao bounds: an evaluation tool for quantitation. *NMR in Biomedicine* **14**, 278–283 (2001).
- Cazzoli, D., Müri, R. M., Schumacher, R., von Arx, S., Chaves, S., Gutbrod, K., Bohlhalter, S., Bauer, D., Vanbellingen, T., Bertschi, M., *et al.* Theta burst stimulation reduces disability during the activities of daily living in spatial neglect. *Brain* **135**, 3426–3439 (2012).
- Cha, H. G. & Kim, M. K. Effects of repetitive transcranial magnetic stimulation on arm function and decreasing unilateral spatial neglect in subacute stroke: a randomized controlled trial. *Clinical rehabilitation* **30**, 649–656 (2016).
- Chandrasekaran, C. Computational principles and models of multisensory integration. *Current opinion in neurobiology* **43**, 25–34 (2017).
- Chang, L., Jiang, C. S. & Ernst, T. Effects of age and sex on brain glutamate and other metabolites. *Magnetic resonance imaging* **27**, 142–145 (2009).
- Chapman, C. A., Perez, Y. & Lacaille, J.-C. Effects of GABAA inhibition on the expression of long-term potentiation in CA1 pyramidal cells are dependent on tetanization parameters. *Hippocampus* **8**, 289–298 (1998).
- Chapman, H. L., Eramudugolla, R., Gavrilescu, M., Strudwick, M. W., Loftus, A., Cunningham, R. & Mattingley, J. B. Neural mechanisms underlying spatial realignment during adaptation to optical wedge prisms. *Neuropsychologia* **48**, 2595–2601 (2010).
- Chaudhuri, O., Gu, L., Darnell, M., Klumpers, D., Bencherif, S. A., Weaver, J. C., Huebsch, N. & Mooney, D. J. Substrate stress relaxation regulates cell spreading. *Nature communications* **6**, 6365 (2015).
- Chechlacz, M., Rotshtein, P., Bickerton, W.-L., Hansen, P. C., Deb, S. & Humphreys, G. W. Separating neural correlates of allocentric and egocentric neglect: distinct cortical sites and common white matter disconnections. *Cognitive Neuropsychology* **27**, 277–303 (2010).

- Chen, P., Hreha, K., Fortis, P., Goedert, K. M. & Barrett, A. M. Functional assessment of spatial neglect: a review of the Catherine Bergego Scale and an introduction of the Kessler Foundation Neglect Assessment Process. *Topics in stroke rehabilitation* **19**, 423–435 (2012).
- Chen, R., Classen, J., Gerloff, C., Celnik, P., Wassermann, E., Hallett, M. & Cohen, L. G. Depression of motor cortex excitability by low-frequency transcranial magnetic stimulation. *Neurology* **48**, 1398–1403 (1997).
- Chen, X., Ye, H., Kuruvilla, R., Ramanan, N., Scangos, K. W., Zhang, C., Johnson, N. M., England, P. M., Shokat, K. M. & Ginty, D. D. A chemical-genetic approach to studying neurotrophin signaling. *Neuron* **46**, 13–21 (2005).
- çiçek, M., Deouell, L. Y. & Knight, R. T. Brain activity during landmark and line bisection tasks. *Frontiers in human neuroscience* **3**, 7 (2009).
- çiçek, M., Gitelman, D., Hurley, R. S., Nobre, A. & Mesulam, M. Anatomical physiology of spatial extinction. *Cerebral Cortex* **17**, 2892–2898 (2007).
- Clower, D. M. & Boussaoud, D. Selective use of perceptual recalibration versus visuomotor skill acquisition. *Journal of Neurophysiology* **84**, 2703–2708 (2000).
- Cohen Kadosh, R., Levy, N., O’Shea, J., Shea, N. & Savulescu, J. The neuroethics of non-invasive brain stimulation. *Current Biology* **22**, R108–R111 (2012).
- Cohen, J. D., Romero, R. D., Servan-Schreiber, D. & Farah, M. J. Mechanisms of spatial attention: The relation of macrostructure to microstructure in parietal neglect. *Journal of Cognitive Neuroscience* **6**, 377–387 (1994).
- Colent, C., Pisella, L., Bernieri, C., Rode, G. & Rossetti, Y. Cognitive bias induced by visuo-motor adaptation to prisms: a simulation of unilateral neglect in normal individuals? *Neuroreport* **11**, 1899–1902 (2000).
- Committeri, G., Pitzalis, S., Galati, G., Patria, F., Pelle, G., Sabatini, U., Castriota-Scanderbeg, A., Piccardi, L., Guariglia, C. & Pizzamiglio, L. Neural bases of personal and extrapersonal neglect in humans. *Brain* **130**, 431–441 (2006).
- Corbetta, M., Kincade, M. J., Lewis, C., Snyder, A. Z. & Sapir, A. Neural basis and recovery of spatial attention deficits in spatial neglect. *Nature neuroscience* **8**, 1603–1610 (2005).
- Corbetta, M., Patel, G. & Shulman, G. L. The reorienting system of the human brain: from environment to theory of mind. *Neuron* **58**, 306–324 (2008).
- Corbetta, M. & Shulman, G. L. Control of goal-directed and stimulus-driven attention in the brain. *Nature reviews neuroscience* **3**, 201–215 (2002).
- Corbetta, M. & Shulman, G. L. Spatial neglect and attention networks. *Annual review of neuroscience* **34**, 569–599 (2011).
- Corbetta, M., Akbudak, E., Conturo, T. E., Snyder, A. Z., Ollinger, J. M., Drury, H. A., Linenweber, M. R., Petersen, S. E., Raichle, M. E., Van Essen, D. C., *et al.* A common network of functional areas for attention and eye movements. *Neuron* **21**, 761–773 (1998).

- Cothros, N., Wong, J. & Gribble, P. Are there distinct neural representations of object and limb dynamics? *Experimental Brain Research* **173**, 689–697 (2006).
- Creutzfeldt, O. D., Fromm, G. H. & Kapp, H. Influence of transcortical dc currents on cortical neuronal activity. *Experimental neurology* **5**, 436–452 (1962).
- Dalmajjer, E. S., Van der Stigchel, S., Nijboer, T. C., Cornelissen, T. H. & Husain, M. CancellationTools: All-in-one software for administration and analysis of cancellation tasks. *Behavior research methods* **47**, 1065–1075 (2015).
- Damasio, A. R., Damasio, H. & Chui, H. C. Neglect following damage to frontal lobe or basal ganglia. *Neuropsychologia* **18**, 123–132 (1980).
- Dambeck, N., Sparing, R., Meister, I. G., Wienemann, M., Weidemann, J., Topper, R. & Boroojerdi, B. Interhemispheric imbalance during visuospatial attention investigated by unilateral and bilateral TMS over human parietal cortices. *Brain research* **1072**, 194–199 (2006).
- Danckert, J., Ferber, S. & Goodale, M. A. Direct effects of prismatic lenses on visuomotor control: an event-related functional MRI study. *European Journal of Neuroscience* **28**, 1696–1704 (2008).
- Datta, A., Baker, J. M., Bikson, M. & Fridriksson, J. Individualized model predicts brain current flow during transcranial direct-current stimulation treatment in responsive stroke patient. *Brain stimulation* **4**, 169–174 (2011).
- Datta, A., Bansal, V., Diaz, J., Patel, J., Reato, D. & Bikson, M. Gyri-precise head model of transcranial direct current stimulation: improved spatial focality using a ring electrode versus conventional rectangular pad. *Brain stimulation* **2**, 201–207 (2009).
- Datta, A., Elwassif, M., Battaglia, F. & Bikson, M. Transcranial current stimulation focality using disc and ring electrode configurations: FEM analysis. *Journal of neural engineering* **5**, 163 (2008).
- Davidson, P. R. & Wolpert, D. M. Widespread access to predictive models in the motor system: a short review. *Journal of Neural Engineering* **2**, S313 (2005).
- Daw, N. D., O’doherly, J. P., Dayan, P., Seymour, B. & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* **441**, 876 (2006).
- De Schotten, M. T., Dell’Acqua, F., Forkel, S. J., Simmons, A., Vergani, F., Murphy, D. G. & Catani, M. A lateralized brain network for visuospatial attention. *Nature neuroscience* **14**, 1245–1246 (2011).
- De Haan, B., Karnath, H.-O. & Driver, J. Mechanisms and anatomy of unilateral extinction after brain injury. *Neuropsychologia* **50**, 1045–1053 (2012).
- Deblieck, C., Thompson, B., Iacoboni, M. & Wu, A. D. Correlation between motor and phosphene thresholds: a transcranial magnetic stimulation study. *Human brain mapping* **29**, 662–670 (2008).
- Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., *et al.* An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* **31**, 968–980 (2006).

- Dewar, R. Adaptation to displaced vision: Variations on the “prismatic-shaping” technique. *Attention, Perception, & Psychophysics* **9**, 155–157 (1971).
- Dhamne, S. C., Kothare, R. S., Yu, C., Hsieh, T.-H., Anastasio, E. M., Oberman, L., Pascual-Leone, A. & Rotenberg, A. A measure of acoustic noise generated from transcranial magnetic stimulation coils. *Brain stimulation* **7**, 432–434 (2014).
- Di Monaco, M., Schintu, S., Dotta, M., Barba, S., Tappero, R. & Gindri, P. Severity of unilateral spatial neglect is an independent predictor of functional outcome after acute inpatient rehabilitation in individuals with right hemispheric stroke. *Archives of physical medicine and rehabilitation* **92**, 1250–1256 (2011).
- Diedrichsen, J. & White, O. *Responsibility Assignment in Redundant Systems Front. Comput. Neurosci. Conference Abstract: Computations, Decisions and Movement. doi: 10.3389/conf.fnins 4* (2010).
- Dmochowski, J. P., Datta, A., Bikson, M., Su, Y. & Parra, L. C. Optimized multi-electrode stimulation increases focality and intensity at target. *Journal of neural engineering* **8**, 046011 (2011).
- Dobek, C. E., Blumberger, D. M., Downar, J., Daskalakis, Z. J. & Vila-Rodriguez, F. Risk of seizures in transcranial magnetic stimulation: a clinical review to inform consent process focused on bupropion. *Neuropsychiatric disease and treatment* **11**, 2975 (2015).
- Doricchi, F., de Schotten, M. T., Tomaiuolo, F. & Bartolomeo, P. White matter (dis)connections and gray matter (dys) functions in visual neglect: gaining insights into the brain networks of spatial awareness. *cortex* **44**, 983–995 (2008).
- Dubljević, V., Saigle, V. & Racine, E. The rising tide of tDCS in the media and academic literature. *Neuron* **82**, 731–736 (2014).
- Duc, A. H., Bays, P. & Husain, M. Eye movements as a probe of attention. *Progress in brain research* **171**, 403–411 (2008).
- Duecker, F. & Sack, A. T. The hybrid model of attentional control: New insights into hemispheric asymmetries inferred from TMS research. *Neuropsychologia* **74**, 21–29 (2015).
- Edwards, D., Cortes, M., Datta, A., Minhas, P., Wassermann, E. M. & Bikson, M. Physiological and modeling evidence for focal transcranial electrical brain stimulation in humans: a basis for high-definition tDCS. *Neuroimage* **74**, 266–275 (2013).
- Egan, M. F., Kojima, M., Callicott, J. H., Goldberg, T. E., Kolachana, B. S., Bertolino, A., Zaitsev, E., Gold, B., Goldman, D., Dean, M., *et al.* The BDNF val66met polymorphism affects activity-dependent secretion of BDNF and human memory and hippocampal function. *Cell* **112**, 257–269 (2003).
- Engel, S. A., Glover, G. H. & Wandell, B. A. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral cortex (New York, NY: 1991)* **7**, 181–192 (1997).

- Epperson, C. N., Haga, K., Mason, G. F., Sellers, E., Gueorguieva, R., Zhang, W., Weiss, E., Rothman, D. L. & Krystal, J. H. Cortical γ -aminobutyric acid levels across the menstrual cycle in healthy women and those with premenstrual dysphoric disorder: A proton magnetic resonance spectroscopy study. *Archives of general psychiatry* **59**, 851–858 (2002).
- Erhardt, J. & Štrac, D. Š. New tools for neuroenhancement—what about neuroethics? *Croatian medical journal* **57**, 392 (2016).
- Ernst, M. O. & Banks, M. S. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415**, 429–433 (2002).
- Ernst, M. O. & Bühlhoff, H. H. Merging the senses into a robust percept. *Trends in cognitive sciences* **8**, 162–169 (2004).
- Ethier, V., Zee, D. S. & Shadmehr, R. Changes in control of saccades during gain adaptation. *Journal of Neuroscience* **28**, 13929–13937 (2008).
- Ethier, V., Zee, D. S. & Shadmehr, R. Spontaneous recovery of motor memory during saccade adaptation. *Journal of Neurophysiology* **99**, 2577–2583 (2008).
- Farnè, A., Rossetti, Y., Toniolo, S. & Làdavas, E. Ameliorating neglect with prism adaptation: visuo-manual and visuo-verbal measures. *Neuropsychologia* **40**, 718–729 (2002).
- Fay, M. P. & Proschan, M. A. Wilcoxon-Mann-Whitney or t-test? On assumptions for hypothesis tests and multiple interpretations of decision rules. *Statistics surveys* **4**, 1 (2010).
- Feigin, V. L., Forouzanfar, M. H., Krishnamurthi, R., Mensah, G. A., Connor, M., Bennett, D. A., Moran, A. E., Sacco, R. L., Anderson, L., Truelsen, T., *et al.* Global and regional burden of stroke during 1990–2010: findings from the Global Burden of Disease Study 2010. *The Lancet* **383**, 245–255 (2014).
- Fernandez-Ruiz, J., Diaz, R., Hall-Haro, C., Vergara, P., Mischner, J., Nunez, L., Drucker-Colin, R., Ochoa, A. & Alonso, M. Normal prism adaptation but reduced after-effect in basal ganglia disorders using a throwing task. *European Journal of Neuroscience* **18**, 689–694 (2003).
- Fernández-Ruiz, J., Hall, C., Vergara, P. & Díaz, R. Prism adaptation in normal aging: slower adaptation rate and larger aftereffect. *Cognitive Brain Research* **9**, 223–226 (2000).
- Fernandez-Ruiz, J., Velásquez-Perez, L., Díaz, R., Drucker-Colín, R., Pérez-González, R., Canales, N., Sánchez-Cruz, G., Martínez-Góngora, E., Medrano, Y., Almaguer-Mederos, L., *et al.* Prism adaptation in spinocerebellar ataxia type 2. *Neuropsychologia* **45**, 2692–2698 (2007).
- Filippini, N., MacIntosh, B. J., Hough, M. G., Goodwin, G. M., Frisoni, G. B., Smith, S. M., Matthews, P. M., Beckmann, C. F. & Mackay, C. E. Distinct patterns of brain activity in young carriers of the APOE- ϵ 4 allele. *Proceedings of the National Academy of Sciences* **106**, 7209–7214 (2009).

- Fink, G. R., Marshall, J. C., Shah, N. J., Weiss, P. H., Halligan, P. W., Grosse-Ruyken, M., Ziemons, K., Zilles, K. & Freund, H.-J. Line bisection judgments implicate right parietal cortex and cerebellum as assessed by fMRI. *Neurology* **54**, 1324–1331 (2000).
- Fisher, R. A. Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. *Biometrika* **10**, 507–521 (1915).
- Fitzgerald, P. B. Transcranial pulsed current stimulation: A new way forward? *Clinical Neurophysiology* **125**, 217–219 (2014).
- Fitzgerald, P. B., Brown, T. L., Daskalakis, Z. J., Chen, R. & Kulkarni, J. Intensity-dependent effects of 1 Hz rTMS on human corticospinal excitability. *Clinical neurophysiology* **113**, 1136–1141 (2002).
- Forkel, S. J., de Schotten, M. T., Kawadler, J. M., Dell’Acqua, F., Danek, A. & Catani, M. The anatomy of fronto-occipital connections from early blunt dissections to contemporary tractography. *Cortex* **56**, 73–84 (2014).
- Formisano, E., Linden, D. E., Di Salle, F., Trojano, L., Esposito, F., Sack, A. T., Grossi, D., Zanella, F. E. & Goebel, R. Tracking the mind’s image in the brain I: time-resolved fMRI during visuospatial mental imagery. *Neuron* **35**, 185–194 (2002).
- Fox, M. D., Corbetta, M., Snyder, A. Z., Vincent, J. L. & Raichle, M. E. Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proceedings of the National Academy of Sciences* **103**, 10046–10051 (2006).
- Fox, P. T. & Raichle, M. E. Focal physiological uncoupling of cerebral blood flow and oxidative metabolism during somatosensory stimulation in human subjects. *Proceedings of the National Academy of Sciences* **83**, 1140–1144 (1986).
- Frank, M. J., Seeberger, L. C. & O’reilly, R. C. By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science* **306**, 1940–1943 (2004).
- Franklin, D. W. & Wolpert, D. M. Computational mechanisms of sensorimotor control. *Neuron* **72**, 425–442 (2011).
- Frassinetti, F., Angeli, V., Meneghello, F., Avanzi, S. & Làdavas, E. Long-lasting amelioration of visuospatial neglect by prism adaptation. *Brain* **125**, 608–623 (2002).
- Frazier, J. A., Chiu, S., Breeze, J. L., Makris, N., Lange, N., Kennedy, D. N., Herbert, M. R., Bent, E. K., Koneru, V. K., Dieterich, M. E., *et al.* Structural brain magnetic resonance imaging of limbic and thalamic volumes in pediatric bipolar disorder. *American Journal of Psychiatry* **162**, 1256–1265 (2005).
- Fritsch, B., Reis, J., Martinowich, K., Schambra, H. M., Ji, Y., Cohen, L. G. & Lu, B. Direct current stimulation promotes BDNF-dependent synaptic plasticity: potential implications for motor learning. *Neuron* **66**, 198–204 (2010).
- Fung, S. H., Roccatagliata, L., Gonzalez, R. G. & Schaefer, P. W. MR diffusion imaging in ischemic stroke. *Neuroimaging clinics of North America* **21**, 345–377 (2011).
- Fusi, S., Drew, P. J. & Abbott, L. F. Cascade models of synaptically stored memories. *Neuron* **45**, 599–611 (2005).

- Galea, J. M., Mallia, E., Rothwell, J. & Diedrichsen, J. The dissociable effects of punishment and reward on motor learning. *Nature neuroscience* **18**, 597–602 (2015).
- Galea, J. M., Vazquez, A., Pasricha, N., Orban de Xivry, J.-J. & Celnik, P. Dissociating the roles of the cerebellum and motor cortex during adaptive learning: the motor cortex retains what the cerebellum learns. *Cerebral cortex* **21**, 1761–1770 (2010).
- Gandiga, P. C., Hummel, F. C. & Cohen, L. G. Transcranial DC stimulation (tDCS): a tool for double-blind sham-controlled clinical studies in brain stimulation. *Clinical neurophysiology* **117**, 845–850 (2006).
- Gao, F., Edden, R. A., Li, M., Puts, N. A., Wang, G., Liu, C., Zhao, B., Wang, H., Bai, X., Zhao, C., *et al.* Edited magnetic resonance spectroscopy detects an age-related decline in brain GABA levels. *Neuroimage* **78**, 75–82 (2013).
- Gartside, I. B. Mechanisms of sustained increases of firing rate of neurones in the rat cerebral cortex after polarization: role of protein synthesis. *Nature* **220**, 383–384 (1968).
- Gauthier, L., Dehaut, F. & Joanette, Y. The bells test: a quantitative and qualitative test for visual neglect. *International journal of clinical neuropsychology* (1989).
- Gerwig, M., Kastrup, O., Meyer, B.-U. & Niehaus, L. Evaluation of cortical excitability by motor and phosphene thresholds in transcranial magnetic stimulation. *Journal of the neurological sciences* **215**, 75–78 (2003).
- Glykys, J. & Mody, I. Activation of GABA A receptors: views from outside the synaptic cleft. *Neuron* **56**, 763–770 (2007).
- Goedert, K. M., LeBlanc, A., Tsai, S.-W. & Barrett, A. M. Asymmetrical effects of adaptation to left-and right-shifting prisms depends on pre-existing attentional biases. *Journal of the International Neuropsychological Society* **16**, 795–804 (2010).
- Good, C. D., Johnsrude, I. S., Ashburner, J., Henson, R. N., Fristen, K. & Frackowiak, R. S. *A voxel-based morphometric study of ageing in 465 normal adult human brains Biomedical Imaging, 2002. 5th IEEE EMBS International Summer School on* (2002), 16–pp.
- Greve, D. N. & Fischl, B. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage* **48**, 63–72 (2009).
- Griffanti, L., Salimi-Khorshidi, G., Beckmann, C. F., Auerbach, E. J., Douaud, G., Sexton, C. E., Zsoldos, E., Ebmeier, K. P., Filippini, N., Mackay, C. E., *et al.* ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *Neuroimage* **95**, 232–247 (2014).
- Guarienti, F., Caumo, W., Shiozawa, P., Cordeiro, Q., Boggio, P. S., Benseñor, I. M., Lotufo, P. A., Bikson, M. & Brunoni, A. R. Reducing Transcranial Direct Current Stimulation-Induced Erythema With Skin Pretreatment: Considerations for Sham-Controlled Clinical Trials. *Neuromodulation: Technology at the Neural Interface* **18**, 261–265 (2015).
- Guigon, E. Interpolation and extrapolation in human behavior and neural networks. *Journal of cognitive neuroscience* **16**, 382–389 (2004).

- Gulani, V., Webb, A., Duncan, I. & Lauterbur, P. Apparent diffusion tensor measurements in myelin-deficient rat spinal cords. *Magnetic Resonance in Medicine* **45**, 191–195 (2001).
- Hadipour-Niktarash, A., Lee, C. K., Desmond, J. E. & Shadmehr, R. Impairment of retention but not acquisition of a visuomotor skill through time-dependent disruption of primary motor cortex. *Journal of Neuroscience* **27**, 13413–13419 (2007).
- Hadjighassem, M., Kamalidehghan, B., Shekarriz, N., Baseerat, A., Molavi, N., Mehrpour, M., Joghataei, M. T., Tondar, M., Ahmadipour, F. & Meng, G. Y. Oral consumption of α -linolenic acid increases serum BDNF levels in healthy adult humans. *Nutrition journal* **14**, 20 (2015).
- Hagmann, P., Jonasson, L., Maeder, P., Thiran, J.-P., Wedeen, V. J. & Meuli, R. Understanding diffusion MR imaging techniques: from scalar diffusion-weighted imaging to diffusion tensor imaging and beyond. *Radiographics* **26**, S205–S223 (2006).
- Haith, A., Jackson, C. P., Miall, R. C. & Vijayakumar, S. *Unifying the sensory and motor components of sensorimotor adaptation Advances in Neural Information Processing Systems* (2009), 593–600.
- Hallett, M. Transcranial magnetic stimulation: a primer. *Neuron* **55**, 187–199 (2007).
- Hamilton, C. R. & Bossom, J. Decay of prism aftereffects. *Journal of experimental psychology* **67**, 148 (1964).
- Hampton, A. N., Bossaerts, P. & O’doherly, J. P. The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *Journal of Neuroscience* **26**, 8360–8367 (2006).
- Hanajima, R., Shadmehr, R., Ohminami, S., Tsutsumi, R., Shiota, Y., Shimizu, T., Tanaka, N., Terao, Y., Tsuji, S., Ugawa, Y., *et al.* Modulation of error-sensitivity during a prism adaptation task in people with cerebellar degeneration. *Journal of neurophysiology* **114**, 2460–2471 (2015).
- Harris, A. D., Puts, N. A. & Edden, R. A. Tissue correction for GABA-edited MRS: Considerations of voxel composition, tissue segmentation, and tissue relaxations. *Journal of Magnetic Resonance Imaging* **42**, 1431–1440 (2015).
- Harris, C. S. Adaptation to displaced vision: visual, motor, or proprioceptive change? *Science* **140**, 812–813 (1963).
- Harris, C. M. & Wolpert, D. M. Signal-dependent noise determines motor planning. *Nature* **394**, 780–784 (1998).
- Hashimoto, Y., Honda, T., Matsumura, K., Nakao, M., Soga, K., Katano, K., Yokota, T., Mizusawa, H., Nagao, S. & Ishikawa, K. Quantitative evaluation of human cerebellum dependent motor learning through prism adaptation of hand-reaching movement. *PloS one* **10**, e0119376 (2015).
- Hatada, Y., Miall, R. & Rossetti, Y. Two waves of a long-lasting aftereffect of prism adaptation measured over 7 days. *Experimental brain research* **169**, 417–426 (2006).

- Hatada, Y., Miall, R. C. & Rossetti, Y. Long lasting aftereffect of a single prism adaptation: directionally biased shift in proprioception and late onset shift of internal egocentric reference frame. *Experimental Brain Research* **174**, 189–198 (2006).
- Hatada, Y., Rossetti, Y. & Miall, R. C. Long-lasting aftereffect of a single prism adaptation: shifts in vision and proprioception are independent. *Experimental brain research* **173**, 415–424 (2006).
- Hay, J. C. & Pick Jr, H. L. Visual and proprioceptive adaptation to optical displacement of the visual stimulus. *Journal of experimental psychology* **71**, 150 (1966).
- Hayes, A. F. *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach* (Guilford Publications, 2017).
- He, B. J., Snyder, A. Z., Vincent, J. L., Epstein, A., Shulman, G. L. & Corbetta, M. Breakdown of functional connectivity in frontoparietal networks underlies behavioral deficits in spatial neglect. *Neuron* **53**, 905–918 (2007).
- He, X., Koo, B.-B. & Killiany, R. J. Edited Magnetic Resonance Spectroscopy Detects an Age-Related Decline in Nonhuman Primate Brain GABA Levels. *BioMed research international* **2016** (2016).
- Hegele, M. & Heuer, H. Adaptation to a direction-dependent visuomotor gain in the young and elderly. *Psychological Research PRPF* **74**, 21 (2010).
- Heilman, K. M. & Van Den Abell, T. Right hemisphere dominance for attention The mechanism underlying hemispheric asymmetries of inattention (neglect). *Neurology* **30**, 327–327 (1980).
- Heilman, K. M., Watson, R. T. & Valenstein, E. Spatial neglect. (2002).
- Heise, K.-F., Zimmerman, M., Hoppe, J., Gerloff, C., Wegscheider, K. & Hummel, F. C. The aging motor system as a model for plastic changes of GABA-mediated intracortical inhibition and their behavioral relevance. *Journal of Neuroscience* **33**, 9039–9049 (2013).
- Held, R. & Hein, A. V. Adaptation of disarranged hand-eye coordination contingent upon re-afferent stimulation. *Perceptual and Motor Skills* **8**, 87–90 (1958).
- Herbet, G., Yordanova, Y. N. & Duffau, H. Left Spatial Neglect Evoked by Electrostimulation of the Right Inferior Fronto-occipital Fasciculus. *Brain Topography*, 1–10 (2017).
- Hermundstad, A. M., Bassett, D. S., Brown, K. S., Aminoff, E. M., Clewett, D., Freeman, S., Frithsen, A., Johnson, A., Tipper, C. M., Miller, M. B., *et al.* Structural foundations of resting-state and task-based functional connectivity in the human brain. *Proceedings of the National Academy of Sciences* **110**, 6169–6174 (2013).
- Herwig, U., Satrapi, P. & Schönfeldt-Lecuona, C. Using the international 10-20 EEG system for positioning of transcranial magnetic stimulation. *Brain topography* **16**, 95–99 (2003).

- Herzfeld, D. J., Pastor, D., Haith, A. M., Rossetti, Y., Shadmehr, R. & O'shea, J. Contributions of the cerebellum and the motor cortex to acquisition and retention of motor memories. *Neuroimage* **98**, 147–158 (2014).
- Heuer, H. & Hegele, M. Age-related variations of visuo-motor adaptation beyond explicit knowledge. *Frontiers in aging neuroscience* **6** (2014).
- Hilgetag, C. C., Théoret, H. & Pascual-Leone, A. Enhanced visual spatial attention ipsilateral to rTMS-induced 'virtual lesions' of human parietal cortex. *Nature neuroscience* **4**, 953–957 (2001).
- Hoge, R. D., Atkinson, J., Gill, B., Crelier, G. R., Marrett, S. & Pike, G. B. Linear coupling between cerebral blood flow and oxygen consumption in activated human cortex. *Proceedings of the National Academy of Sciences* **96**, 9403–9408 (1999).
- Honey, C., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.-P., Meuli, R. & Hagmann, P. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences* **106**, 2035–2040 (2009).
- Hsu, M., Bhatt, M., Adolphs, R., Tranel, D. & Camerer, C. F. Neural systems responding to degrees of uncertainty in human decision-making. *Science* **310**, 1680–1683 (2005).
- Huang, H. J. & Ahmed, A. A. Older adults learn less, but still reduce metabolic cost, during motor adaptation. *Journal of neurophysiology* **111**, 135–144 (2014).
- Huang, Y., Liu, A. A., Lafon, B., Friedman, D., Dayan, M., Wang, X., Bikson, M., Doyle, W. K., Devinsky, O. & Parra, L. C. Measurements and models of electric fields in the in vivo human brain during transcranial electric stimulation. *Elife* **6**, e18834 (2017).
- Huberdeau, D. M., Krakauer, J. W. & Haith, A. M. Dual-process decomposition in human sensorimotor adaptation. *Current opinion in neurobiology* **33**, 71–77 (2015).
- Hummel, F., Celnik, P., Giraux, P., Floel, A., Wu, W.-H., Gerloff, C. & Cohen, L. G. Effects of non-invasive cortical stimulation on skilled motor function in chronic stroke. *Brain* **128**, 490–499 (2005).
- Hunter, T., Sacco, P., Nitsche, M. A. & Turner, D. L. Modulation of internal model formation during force field-induced motor learning by anodal transcranial direct current stimulation of primary motor cortex. *The Journal of physiology* **587**, 2949–2961 (2009).
- Husain, M. & Kennard, C. Visual neglect associated with frontal lobe infarction. *Journal of neurology* **243**, 652–657 (1996).
- Husain, M., Mannan, S., Hodgson, T., Wojciulik, E., Driver, J. & Kennard, C. Impaired spatial working memory across saccades contributes to abnormal search in parietal neglect. *Brain* **124**, 941–952 (2001).
- Husain, M., Shapiro, K., Martin, J. & Kennard, C. Abnormal temporal dynamics of visual attention in spatial neglect patients. *Nature* **385**, 154 (1997).

- Inoue, M., Uchimura, M., Karibe, A., O'Shea, J., Rossetti, Y. & Kitazawa, S. Three timescales in prism adaptation. *Journal of neurophysiology* **113**, 328–338 (2015).
- Inoue, M., Uchimura, M. & Kitazawa, S. Error signals in motor cortices drive adaptation in reaching. *Neuron* **90**, 1114–1126 (2016).
- Ip, I. B., Berrington, A., Hess, A. T., Parker, A. J., Emir, U. E. & Bridge, H. Combined fMRI-MRS acquires simultaneous glutamate and BOLD-fMRI signals in the human brain. *NeuroImage* **155**, 113–119 (2017).
- Jacquin-Courtois, S. Hemi-spatial neglect rehabilitation using non-invasive brain stimulation: or how to modulate the disconnection syndrome? *Annals of physical and rehabilitation medicine* **58**, 251–258 (2015).
- Jacquin-Courtois, S., Rode, G., Pavani, F., O'Shea, J., Giard, M., Boisson, D. & Rossetti, Y. Effect of prism adaptation on left dichotic listening deficit in neglect patients: glasses to hear better? *Brain* **133**, 895–908 (2010).
- Jacquin-Courtois, S., O'shea, J., Luauté, J., Pisella, L., Revol, P., Mizuno, K., Rode, G. & Rossetti, Y. Rehabilitation of spatial neglect by prism adaptation: a peculiar expansion of sensorimotor after-effects to spatial cognition. *Neuroscience & Biobehavioral Reviews* **37**, 594–609 (2013).
- Jacquin-Courtois, S., Rode, G., Pisella, L., Boisson, D. & Rossetti, Y. Wheel-chair driving improvement following visuo-manual prism adaptation. *Cortex* **44**, 90–96 (2008).
- Jehkonen, M., Laihosalo, M. & Kettunen, J. Impact of neglect on functional outcome after stroke—a review of methodological issues and recent research findings. *Restorative neurology and neuroscience* **24**, 209–215 (2006).
- Jenkinson, M., Bannister, P., Brady, M. & Smith, S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* **17**, 825–841 (2002).
- Jenkinson, M., Beckmann, C. F., Behrens, T. E., Woolrich, M. W. & Smith, S. M. Fsl. *Neuroimage* **62**, 782–790 (2012).
- Jenkinson, M. & Smith, S. A global optimisation method for robust affine registration of brain images. *Medical image analysis* **5**, 143–156 (2001).
- Jenkinson, M., Pechaud, M., Smith, S., *et al.* BET2: MR-based estimation of brain, skull and scalp surfaces *Eleventh annual meeting of the organization for human brain mapping* **17** (2005), 167.
- Jewell, G. & McCourt, M. E. Pseudoneglect: a review and meta-analysis of performance factors in line bisection tasks. *Neuropsychologia* **38**, 93–110 (2000).
- Jin, P. Toward a reconceptualization of the law of initial value. *Psychological bulletin* **111**, 176 (1992).
- Jocham, G., Hunt, L. T., Near, J. & Behrens, T. E. A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex. *Nature neuroscience* **15**, 960–961 (2012).

- Johansen-Berg, H. & Behrens, T. E. *Diffusion MRI: from quantitative measurement to in vivo neuroanatomy* (Academic Press, 2013).
- Joiner, W. M. & Smith, M. A. Long-term retention explained by a model of short-term learning in the adaptive control of reaching. *Journal of neurophysiology* **100**, 2948–2955 (2008).
- Joundi, R. A., Lopez-Alonso, V., Lago, A., Brittain, J.-S., Fernandez-del-Olmo, M., Gomez-Garre, P., Mir, P., Jenkinson, N., Cheeran, B. & Brown, P. The effect of BDNF val66met polymorphism on visuomotor adaptation. *Experimental brain research* **223**, 43–50 (2012).
- Juchem, C. & de Graaf, R. A. B0 magnetic field homogeneity and shimming for in vivo magnetic resonance spectroscopy. *Analytical biochemistry* **529**, 17–29 (2017).
- Kaiser, L. G., Schuff, N., Cashdollar, N. & Weiner, M. W. Age-related glutamate and glutamine concentration changes in normal human brain: 1 H MR spectroscopy study at 4 T. *Neurobiology of aging* **26**, 665–672 (2005).
- Kang, E. K. & Paik, N.-J. Effect of a tDCS electrode montage on implicit motor sequence learning in healthy subjects. *Experimental & translational stroke medicine* **3**, 4 (2011).
- Karnath, H.-O. Spatial limitation of eye movements during ocular exploration of simple line drawings in neglect syndrome. *Cortex* **30**, 319–330 (1994).
- Karnath, H.-O., Ferber, S. & Himmelbach, M. Spatial awareness is a function of the temporal not the posterior parietal lobe. *Nature* **411**, 950–953 (2001).
- Karnath, H.-O. & Fetter, M. Ocular space exploration in the dark and its relation to subjective and objective body orientation in neglect patients with parietal lesions. *Neuropsychologia* **33**, 371–377 (1995).
- Karnath, H.-O., Fruhmann Berger, M., Küker, W. & Rorden, C. The anatomy of spatial neglect based on voxelwise statistical analysis: a study of 140 patients. *Cerebral Cortex* **14**, 1164–1172 (2004).
- Karnath, H.-O., Rennig, J., Johannsen, L. & Rorden, C. The anatomy underlying acute versus chronic spatial neglect: a longitudinal study. *Brain* **134**, 903–912 (2010).
- Kawato, M. Internal models for motor control and trajectory planning. *Current opinion in neurobiology* **9**, 718–727 (1999).
- Kayser, C. & Shams, L. Multisensory causal inference in the brain. *PLoS biology* **13**, e1002075 (2015).
- Kessler, S. K., Turkeltaub, P. E., Benson, J. G. & Hamilton, R. H. Differences in the experience of active and sham transcranial direct current stimulation. *Brain stimulation* **5**, 155–162 (2012).
- Kim, S., Stephenson, M. C., Morris, P. G. & Jackson, S. R. tDCS-induced alterations in GABA concentration within primary motor cortex predict motor learning and motor memory: A 7T magnetic resonance spectroscopy study. *Neuroimage* **99**, 237–243 (2014).

- Kim, S., Ogawa, K., Lv, J., Schweighofer, N. & Imamizu, H. Neural substrates related to motor memory with multiple timescales in sensorimotor adaptation. *PLoS biology* **13**, e1002312 (2015).
- King, B. R., Fogel, S. M., Albouy, G. & Doyon, J. Neural correlates of the age-related changes in motor sequence learning and motor adaptation in older adults. *Frontiers in human neuroscience* **7** (2013).
- Kinsbourne, M. Hemi-neglect and hemisphere rivalry. *Advances in neurology* **18**, 41–49 (1976).
- Kitazawa, S., Kimura, T. & Uka, T. Prism adaptation of reaching movements: specificity for the velocity of reaching. *Journal of Neuroscience* **17**, 1481–1492 (1997).
- Kitazawa, S., Kohno, T. & Uka, T. Effects of delayed visual information on the rate and amount of prism adaptation in the human. *Journal of Neuroscience* **15**, 7644–7652 (1995).
- Kleim, J. A., Chan, S., Pringle, E., Schallert, K., Procaccio, V., Jimenez, R. & Cramer, S. C. BDNF val66met polymorphism is associated with modified experience-dependent plasticity in human motor cortex. *Nature neuroscience* **9**, 735–737 (2006).
- Kluzik, J., Diedrichsen, J., Shadmehr, R. & Bastian, A. J. Reach adaptation: what determines whether we learn an internal model of the tool or adapt the model of our arm? *Journal of neurophysiology* **100**, 1455–1464 (2008).
- Ko, M.-H., Han, S.-H., Park, S.-H., Seo, J.-H. & Kim, Y.-H. Improvement of visual scanning after DC brain polarization of parietal cortex in stroke patients with spatial neglect. *Neuroscience letters* **448**, 171–174 (2008).
- Koch, G., Bonni, S., Giacobbe, V., Bucci, G., Basile, B., Lupo, F., Versace, V., Bozzali, M. & Caltagirone, C. Theta-burst stimulation of the left hemisphere accelerates recovery of hemispatial neglect. *Neurology* **78**, 24–30 (2012).
- Koch, G., Oliveri, M., Cheeran, B., Ruge, D., Gerfo, E. L., Salerno, S., Torriero, S., Marconi, B., Mori, F., Driver, J., *et al.* Hyperexcitability of parietal-motor functional connections in the intact left-hemisphere of patients with neglect. *Brain* **131**, 3147–3155 (2008).
- Kolasinski, J., Logan, J. P., Hinson, E. L., Manners, D., Zand, A. P. D., Makin, T. R., Emir, U. E. & Stagg, C. J. A Mechanistic Link from GABA to Cortical Architecture and Perception. *Current Biology* (2017).
- Körding, K. P., Tenenbaum, J. B. & Shadmehr, R. The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature neuroscience* **10**, 779 (2007).
- Körding, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. *Nature* **427**, 244 (2004).
- Körding, K. P. & Wolpert, D. M. Bayesian decision theory in sensorimotor control. *Trends in cognitive sciences* **10**, 319–326 (2006).

- Körding, K. P. & Wolpert, D. M. The loss function of sensorimotor learning. *Proceedings of the National Academy of Sciences of the United States of America* **101**, 9839–9842 (2004).
- Kornheiser, A. S. Adaptation to laterally displaced vision: A review. *Psychological bulletin* **83**, 783 (1976).
- Kosslyn, S. M., Pascual-Leone, A., Felician, O., Camposano, S., Keenan, J., Ganis, G., Sukel, K., Alpert, N., *et al.* The role of area 17 in visual imagery: convergent evidence from PET and rTMS. *Science* **284**, 167–170 (1999).
- Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A. & Poeppel, D. Neuroscience needs behavior: correcting a reductionist Bias. *Neuron* **93**, 480–490 (2017).
- Krakauer, J. W., Pine, Z. M., Ghilardi, M.-F. & Ghez, C. Learning of visuomotor transformations for vectorial planning of reaching trajectories. *Journal of Neuroscience* **20**, 8916–8924 (2000).
- Krakauer, J. W. & Shadmehr, R. Consolidation of motor memory. *Trends in neurosciences* **29**, 58–64 (2006).
- Krause, B., Márquez-Ruiz, J. & Kadosh, R. C. The effect of transcranial direct current stimulation: a role for cortical excitation/inhibition balance? *Frontiers in human neuroscience* **7** (2013).
- Kreis, R. The trouble with quality filtering based on relative Cramér-Rao lower bounds. *Magnetic resonance in medicine* **75**, 15–18 (2016).
- Kuo, H.-I., Bikson, M., Datta, A., Minhas, P., Paulus, W., Kuo, M.-F. & Nitsche, M. A. Comparing cortical plasticity induced by conventional and high-definition 4 × 1 ring tDCS: a neurophysiological study. *Brain stimulation* **6**, 644–648 (2013).
- Küper, M., Wünnemann, M. J., Thürling, M., Stefanescu, R. M., Maderwald, S., Elles, H. G., Göricke, S., Ladd, M. E. & Timmann, D. Activation of the cerebellar cortex and the dentate nucleus in a prism adaptation fMRI study. *Human brain mapping* **35**, 1574–1586 (2014).
- Kurata, K. & Hoshi, E. Reacquisition deficits in prism adaptation after muscimol microinjection into the ventral premotor cortex of monkeys. *Journal of neurophysiology* **81**, 1927–1938 (1999).
- Lackner, J. R. & Dizio, P. Rapid adaptation to Coriolis force perturbations of arm trajectory. *Journal of neurophysiology* **72**, 299–313 (1994).
- Làdavas, E., Bonifazi, S., Catena, L. & Serino, A. Neglect rehabilitation by prism adaptation: different procedures have different impacts. *Neuropsychologia* **49**, 1136–1145 (2011).
- Lalazar, H. & Vaadia, E. Neural basis of sensorimotor learning: modifying internal models. *Current opinion in neurobiology* **18**, 573–581 (2008).

- Landi, S. M., Baguear, F. & Della-Maggiore, V. One week of motor adaptation induces structural changes in primary motor cortex that predict long-term memory one year later. *Journal of Neuroscience* **31**, 11808–11813 (2011).
- Lang, N., Harms, J., Weyh, T., Lemon, R. N., Paulus, W., Rothwell, J. C. & Siebner, H. R. Stimulus intensity and coil characteristics influence the efficacy of rTMS to suppress cortical excitability. *Clinical Neurophysiology* **117**, 2292–2301 (2006).
- Lang, N., Siebner, H. R., Ernst, D., Nitsche, M. A., Paulus, W., Lemon, R. N. & Rothwell, J. C. Preconditioning with transcranial direct current stimulation sensitizes the motor cortex to rapid-rate transcranial magnetic stimulation and controls the direction of after-effects. *Biological psychiatry* **56**, 634–639 (2004).
- Lange, I., Kasanova, Z., Goossens, L., Leibold, N., De Zeeuw, C. I., van Amelsvoort, T. & Schruers, K. The anatomy of fear learning in the cerebellum: A systematic meta-analysis. *Neuroscience & Biobehavioral Reviews* **59**, 83–91 (2015).
- Langhorne, P., Bernhardt, J. & Kwakkel, G. Stroke rehabilitation. *The Lancet* **377**, 1693–1702 (2011).
- Le Bihan, D. & Johansen-Berg, H. Diffusion MRI at 25: exploring brain tissue structure and function. *Neuroimage* **61**, 324–341 (2012).
- Le, A., Stojanoski, B. B., Khan, S., Keough, M. & Niemeier, M. A toggle switch of visual awareness? *Cortex* **64**, 169–178 (2015).
- Le, T. H., Pardo, J. V. & Hu, X. 4 T-fMRI study of nonspatial shifting of selective attention: cerebellar and parietal contributions. *Journal of Neurophysiology* **79**, 1535–1548 (1998).
- Lee, J.-Y. & Schweighofer, N. Dual adaptation supports a parallel architecture of motor memory. *Journal of Neuroscience* **29**, 10396–10404 (2009).
- Leibovitch, F., Black, S., Caldwell, C., Ebert, P., Ehrlich, L. & Szalai, J. Brain-behavior correlations in hemispatial neglect using CT and SPECT The Sunnybrook Stroke Study. *Neurology* **50**, 901–908 (1998).
- Leow, L.-A., Marinovic, W., Chye, L., Carroll, T. J. & Riek, S. The primary motor cortex is critical for the retention of implicit sensorimotor adaptation. *bioRxiv*, 096404 (2016).
- Levin, O., Fujiyama, H., Boisgontier, M. P., Swinnen, S. P. & Summers, J. J. Aging and motor inhibition: a converging perspective provided by brain stimulation and imaging approaches. *Neuroscience & Biobehavioral Reviews* **43**, 100–117 (2014).
- Levkovitz, Y., Avignone, E., Groner, Y. & Segal, M. Upregulation of GABA neurotransmission suppresses hippocampal excitability and prevents long-term potentiation in transgenic superoxide dismutase-overexpressing mice. *Journal of Neuroscience* **19**, 10977–10984 (1999).
- Li, C.-S. R., Padoa-Schioppa, C. & Bizzi, E. Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* **30**, 593–607 (2001).

- Liebetanz, D., Nitsche, M. A., Tergau, F. & Paulus, W. Pharmacological approach to the mechanisms of transcranial DC-stimulation-induced after-effects of human motor cortex excitability. *Brain* **125**, 2238–2247 (2002).
- Lim, J. Y., Kang, E. K. & Paik, N.-J. Repetitive transcranial magnetic stimulation for hemispatial neglect in patients after stroke: an open-label pilot study. *Journal of rehabilitation medicine* **42**, 447–452 (2010).
- Loftus, A. M., Vijayakumar, N. & Nicholls, M. E. Prism adaptation overcomes pseudoneglect for the greyscales task. *Cortex* **45**, 537–543 (2009).
- Luauté, J., Schwartz, S., Rossetti, Y., Spiridon, M., Rode, G., Boisson, D. & Vuilleumier, P. Dynamic changes in brain activity during prism adaptation. *Journal of Neuroscience* **29**, 169–178 (2009).
- Luauté, J., Michel, C., Rode, G., Pisella, L., Jacquin-Courtois, S., Costes, N., Cotton, F., Le Bars, D., Boisson, D., Halligan, P., *et al.* Functional anatomy of the therapeutic effects of prism adaptation on left neglect. *Neurology* **66**, 1859–1867 (2006).
- Lumer, E. D., Friston, K. J. & Rees, G. Neural correlates of perceptual rivalry in the human brain. *Science* **280**, 1930–1934 (1998).
- Lunghi, C., Emir, U. E., Morrone, M. C. & Bridge, H. Short-term monocular deprivation alters GABA in the adult human visual cortex. *Current Biology* **25**, 1496–1501 (2015).
- Lunven, M., Thiebaut de Schotten, M., Bourlon, C., Duret, C., Migliaccio, R., Rode, G. & Bartolomeo, P. White matter lesional predictors of chronic visual neglect: a longitudinal study. *Brain* **138**, 746–760 (2015).
- Lyle, R. C. A performance test for assessment of upper limb function in physical rehabilitation treatment and research. *International Journal of Rehabilitation Research* **4**, 483–492 (1981).
- Macaluso, E. & Patria, F. Spatial re-orienting of visual attention along the horizontal or the vertical axis. *Experimental brain research* **180**, 23–34 (2007).
- Macey, P. M., Macey, K. E., Kumar, R. & Harper, R. M. A method for removal of global effects from fMRI time series. *Neuroimage* **22**, 360–366 (2004).
- MacFarland, T. W. & Yates, J. M. *Introduction to Nonparametric Statistics for the Biological Sciences Using R* 103–132 (Springer, 2016).
- Maeda, F., Keenan, J. P., Tormos, J. M., Topka, H. & Pascual-Leone, A. Modulation of corticospinal excitability by repetitive transcranial magnetic stimulation. *Clinical Neurophysiology* **111**, 800–805 (2000).
- Maffei, A., Charrier, C., Caiati, M. D., Barberis, A., Mahadevan, V., Woodin, M. A. & Tyagarajan, S. K. Emerging Mechanisms Underlying Dynamics of GABAergic Synapses. *Journal of Neuroscience* **37**, 10792–10799 (2017).
- Magistretti, P. J., Pellerin, L., Rothman, D. L. & Shulman, R. G. Energy on demand. *Science* **283**, 496–497 (1999).

Mann, H. B. & Whitney, D. R. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, 50–60 (1947).

Markett, S., Reuter, M., Montag, C., Voigt, G., Lachmann, B., Rudolf, S., Elger, C. E. & Weber, B. Assessing the function of the fronto-parietal attention network: insights from resting-state fMRI and the attentional network test. *Human brain mapping* **35**, 1700–1709 (2014).

Marr, D. A computational investigation into the human representation and processing of visual information. *WH San Francisco: Freeman and Company* **1** (1982).

Mars, R. B., Sallet, J., Schüffelen, U., Jbabdi, S., Toni, I. & Rushworth, M. F. Connectivity-based subdivisions of the human right “temporoparietal junction area”: evidence for different areas participating in different cortical networks. *Cerebral cortex* **22**, 1894–1903 (2011).

Mars, R. B., Jbabdi, S., Sallet, J., O’Reilly, J. X., Croxson, P. L., Olivier, E., Noonan, M. P., Bergmann, C., Mitchell, A. S., Baxter, M. G., *et al.* Diffusion-weighted imaging tractography-based parcellation of the human parietal cortex and comparison with human and macaque resting-state functional connectivity. *Journal of Neuroscience* **31**, 4087–4100 (2011).

Martin, L. J., Zurek, A. A., MacDonald, J. F., Roder, J. C., Jackson, M. F. & Orser, B. A. $\alpha 5$ GABAA receptor activity sets the threshold for long-term potentiation and constrains hippocampus-dependent memory. *Journal of Neuroscience* **30**, 5269–5282 (2010).

Martin, T., Keating, J., Goodkin, H., Bastian, A. & Thach, W. Throwing while looking through prisms: I. Focal olivocerebellar lesions impair adaptation. *Brain* **119**, 1183–1198 (1996).

Martín-Arévalo, E., Laube, I., Koun, E., Farnè, A., Reilly, K. T. & Pisella, L. Prism adaptation alters electrophysiological markers of attentional processes in the healthy brain. *Journal of Neuroscience* **36**, 1019–1030 (2016).

Mattingley, J. B., Berberovic, N., Corben, L., Slavin, M. J., Nicholls, M. E. & Bradshaw, J. L. The greyscales task: a perceptual measure of attentional bias following unilateral hemispheric damage. *Neuropsychologia* **42**, 387–394 (2004).

Mazzoni, P., Hristova, A. & Krakauer, J. W. Why don’t we move faster? Parkinson’s disease, movement vigor, and implicit motivation. *Journal of neuroscience* **27**, 7105–7116 (2007).

Mazzoni, P. & Krakauer, J. W. An implicit plan overrides an explicit strategy during visuomotor adaptation. *Journal of neuroscience* **26**, 3642–3645 (2006).

McConnell, K. A., Nahas, Z., Shastri, A., Lorberbaum, J. P., Kozel, F. A., Bohning, D. E. & George, M. S. The transcranial magnetic stimulation motor threshold depends on the distance from coil to underlying cortex: a replication in healthy adults comparing two methods of assessing the distance to cortex. *Biological psychiatry* **49**, 454–459 (2001).

- McDougle, S. D., Bond, K. M. & Taylor, J. A. Explicit and implicit processes constitute the fast and slow processes of sensorimotor learning. *Journal of Neuroscience* **35**, 9568–9579 (2015).
- McHughen, S. A., Pearson-Fuhrhop, K., Ngo, V. K. & Cramer, S. C. Intense training overcomes effects of the Val66Met BDNF polymorphism on short-term plasticity. *Experimental brain research* **213**, 415 (2011).
- McKeown, M. J., Makeig, S., Brown, G. G., Jung, T.-P., Kindermann, S. S., Bell, A. J. & Sejnowski, T. J. *Analysis of fMRI data by blind separation into independent spatial components* tech. rep. (Naval Health Research Center San Diego CA, 1997).
- Meister, I., Wienemann, M., Buelte, D., Grünewald, C., Sparing, R., Dambeck, N. & Boroojerdi, B. Hemiextinction induced by transcranial magnetic stimulation over the right temporo-parietal junction. *Neuroscience* **142**, 119–123 (2006).
- Mesulam, M. *et al.* A cortical network for directed attention and unilateral neglect. *Annals of neurology* **10**, 309–325 (1981).
- Miall, R. C. & Wolpert, D. M. Forward models for physiological motor control. *Neural networks* **9**, 1265–1279 (1996).
- Michel, C., Pisella, L., Halligan, P. W., Luauté, J., Rode, G., Boisson, D. & Rossetti, Y. Simulating unilateral neglect in normals using prism adaptation: implications for theory. *Neuropsychologia* **41**, 25–39 (2003).
- Michel, C., Pisella, L., Prablanc, C., Rode, G. & Rossetti, Y. Enhancing visuomotor adaptation by reducing error signals: single-step (aware) versus multiple-step (unaware) exposure to wedge prisms. *Journal of Cognitive Neuroscience* **19**, 341–350 (2007).
- Milner, A. D. & McIntosh, R. D. The neurological basis of visual neglect. *Current opinion in neurology* **18**, 748–753 (2005).
- Minarik, T., Berger, B., Althaus, L., Bader, V., Biebl, B., Brotzeller, F., Fusban, T., Hegemann, J., Jesteadt, L., Kalweit, L., *et al.* The importance of sample size for reproducibility of tDCS effects. *Frontiers in human neuroscience* **10** (2016).
- Minichiello, L. TrkB signalling pathways in LTP and learning. *Nature reviews. Neuroscience* **10**, 850 (2009).
- Miranda, P. C., Lomarev, M. & Hallett, M. Modeling the current distribution during transcranial direct current stimulation. *Clinical neurophysiology* **117**, 1623–1629 (2006).
- Mizuno, K., Tsuji, T., Takebayashi, T., Fujiwara, T., Hase, K. & Liu, M. Prism adaptation therapy enhances rehabilitation of stroke patients with unilateral spatial neglect: a randomized, controlled trial. *Neurorehabilitation and neural repair* **25**, 711–720 (2011).
- Mody, I. Distinguishing between GABAA receptors responsible for tonic and phasic conductances. *Neurochemical research* **26**, 907–913 (2001).

- Monte-Silva, K., Kuo, M.-F., Hessenthaler, S., Fresnoza, S., Liebetanz, D., Paulus, W. & Nitsche, M. A. Induction of late LTP-like plasticity in the human motor cortex by repeated non-invasive brain stimulation. *Brain stimulation* **6**, 424–432 (2013).
- Morasso, P. Spatial control of arm movements. *Experimental brain research* **42**, 223–227 (1981).
- Mori, S., Wakana, S., Van Zijl, P. C. & Nagae-Poetscher, L. *MRI atlas of human white matter* (Elsevier, 2005).
- Mort, D. J., Malhotra, P., Mannan, S. K., Rorden, C., Pambakian, A., Kennard, C. & Husain, M. The anatomy of visual neglect. *Brain* **126**, 1986–1997 (2003).
- Moseley, M., Kucharczyk, J., Mintorovitch, J., Cohen, Y., Kurhanewicz, J., Derugin, N., Asgari, H. & Norman, D. Diffusion-weighted MR imaging of acute stroke: correlation with T2-weighted and magnetic susceptibility-enhanced MR imaging in cats. *American Journal of Neuroradiology* **11**, 423–429 (1990).
- Murphy, K., Birn, R. M. & Bandettini, P. A. Resting-state fMRI confounds and cleanup. *Neuroimage* **80**, 349–359 (2013).
- Natt, O., Bezkorovaynyy, V., Michaelis, T. & Frahm, J. Use of phased array coils for a determination of absolute metabolite concentrations. *Magnetic resonance in medicine* **53**, 3–8 (2005).
- Nemanich, S. T. & Earhart, G. M. How do age and nature of the motor task influence visuomotor adaptation? *Gait & posture* **42**, 564–568 (2015).
- Newport, R. & Jackson, S. Posterior parietal cortex and the dissociable components of prism adaptation. *Neuropsychologia* **44**, 2757–2765 (2006).
- Newport, R., Brown, L., Husain, M., Mort, D. & Jackson, S. R. The role of the posterior parietal lobe in prism adaptation: Failure to adapt to optical prisms in a patient with bilateral damage to posterior parietal cortex. *Cortex* **42**, 720–729 (2006).
- Newport, R. & Schenk, T. Prisms and neglect: what have we learned? *Neuropsychologia* **50**, 1080–1091 (2012).
- Nicholls, M. E., Bradshaw, J. L. & Mattingley, J. B. Free-viewing perceptual asymmetries for the judgement of brightness, numerosity and size. *Neuropsychologia* **37**, 307–314 (1999).
- Nitsche, M., Fricke, K., Henschke, U., Schlitterlau, A., Liebetanz, D., Lang, N., Henning, S., Tergau, F. & Paulus, W. Pharmacological modulation of cortical excitability shifts induced by transcranial direct current stimulation in humans. *The Journal of physiology* **553**, 293–301 (2003).
- Nitsche, M., Niehaus, L., Hoffmann, K., Hengst, S., Liebetanz, D., Paulus, W. & Meyer, B.-U. MRI study of human brain exposed to weak direct current stimulation of the frontal cortex. *Clinical Neurophysiology* **115**, 2419–2423 (2004).
- Nitsche, M. A., Liebetanz, D., Schlitterlau, A., Henschke, U., Fricke, K., Frommann, K., Lang, N., Henning, S., Paulus, W. & Tergau, F. GABAergic modulation of DC stimulation-induced motor cortex excitability shifts in humans. *European Journal of Neuroscience* **19**, 2720–2726 (2004).

- Nitsche, M. A. & Paulus, W. Excitability changes induced in the human motor cortex by weak transcranial direct current stimulation. *The Journal of physiology* **527**, 633–639 (2000).
- Nitsche, M. A. & Paulus, W. Sustained excitability elevations induced by transcranial DC motor cortex stimulation in humans. *Neurology* **57**, 1899–1901 (2001).
- Nitsche, M. A., Seeber, A., Frommann, K., Klein, C. C., Rochford, C., Nitsche, M. S., Fricke, K., Liebetanz, D., Lang, N., Antal, A., *et al.* Modulating parameters of excitability during and after transcranial direct current stimulation of the human motor cortex. *The Journal of physiology* **568**, 291–303 (2005).
- Nys, G., De Haan, E., Kunneman, A., De Kort, P. & Dijkerman, H. Acute neglect rehabilitation using repetitive prism adaptation: a randomized placebo-controlled trial. *Restorative neurology and neuroscience* **26**, 1–12 (2008).
- O'Brien, J. S. & Sampson, E. L. Lipid composition of the normal human brain: gray matter, white matter, and myelin. *Journal of lipid research* **6**, 537–544 (1965).
- O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H. & Dolan, R. J. Temporal difference models and reward-related learning in the human brain. *Neuron* **38**, 329–337 (2003).
- O'doherty, J. P., Hampton, A. & Kim, H. Model-based fMRI and its application to reward learning and decision making. *Annals of the New York Academy of sciences* **1104**, 35–53 (2007).
- Ogawa, S., Tank, D. W., Menon, R., Ellermann, J. M., Kim, S. G., Merkle, H. & Ugurbil, K. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences* **89**, 5951–5955 (1992).
- Okamoto, M., Dan, H., Sakamoto, K., Takeo, K., Shimizu, K., Kohno, S., Oda, I., Isobe, S., Suzuki, T., Kohyama, K., *et al.* Three-dimensional probabilistic anatomical cranio-cerebral correlation via the international 10–20 system oriented for transcranial functional brain mapping. *Neuroimage* **21**, 99–111 (2004).
- Oldfield, R. C. The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia* **9**, 97–113 (1971).
- Oldham, P. A note on the analysis of repeated measurements of the same subjects. *Journal of chronic diseases* **15**, 969–977 (1962).
- Oliveri, M., Bisiach, E., Brighina, F., Piazza, A., La Bua, V., Buffa, D. & Fierro, B. rTMS of the unaffected hemisphere transiently reduces contralesional visuospatial hemineglect. *Neurology* **57**, 1338–1340 (2001).
- Oliviero, A., Profice, P., Tonali, P., Pilato, F., Saturno, E., Dileone, M., Ranieri, F. & Di Lazzaro, V. Effects of aging on motor cortex excitability. *Neuroscience research* **55**, 74–77 (2006).
- Orban de Xivry, J.-J., Marko, M. K., Pekny, S. E., Pastor, D., Izawa, J., Celnik, P. & Shadmehr, R. Stimulation of the human motor cortex alters generalization patterns of motor learning. *Journal of Neuroscience* **31**, 7102–7110 (2011).

- O'Reilly, J. X., Mesulam, M. M. & Nobre, A. C. The cerebellum predicts the timing of perceptual events. *Journal of Neuroscience* **28**, 2252–2260 (2008).
- O'Shea, J., Boudrias, M.-H., Stagg, C. J., Bachtiar, V., Kischka, U., Blicher, J. U. & Johansen-Berg, H. Predicting behavioural response to TDCS in chronic motor stroke. *Neuroimage* **85**, 924–933 (2014).
- O'Shea, J., Gaveau, V., Kandel, M., Koga, K., Susami, K., Prablanc, C. & Rossetti, Y. Kinematic markers dissociate error correction from sensorimotor realignment during prism adaptation. *Neuropsychologia* **55**, 15–24 (2014).
- O'Shea, J., Johansen-Berg, H., Trief, D., Göbel, S. & Rushworth, M. F. Functionally specific reorganization in human premotor cortex. *Neuron* **54**, 479–490 (2007).
- O'Shea, J., Revol, P., Cousijn, H., Near, J., Petitot, P., Jacquin-Courtois, S., Johansen-Berg, H., Rode, G. & Rossetti, Y. Induced sensorimotor cortex plasticity remediates chronic treatment-resistant visual neglect. *eLife* **6** (2017).
- öz, G. & Tkáč, I. Short-echo, single-shot, full-intensity proton magnetic resonance spectroscopy for neurochemical profiling at 4 T: Validation in the cerebellum and brainstem. *Magnetic resonance in medicine* **65**, 901–910 (2011).
- Panouillères, M. T., Joundi, R. A., Brittain, J.-S. & Jenkinson, N. Reversing motor adaptation deficits in the ageing brain using non-invasive stimulation. *The Journal of physiology* **593**, 3645–3655 (2015).
- Parazzini, M., Fiocchi, S., Rossi, E., Paglialonga, A. & Ravazzani, P. Transcranial direct current stimulation: estimation of the electric field and of the current density in an anatomical human head model. *IEEE Transactions on Biomedical Engineering* **58**, 1773–1780 (2011).
- Parks, E. L. & Madden, D. J. Brain connectivity and visual attention. *Brain connectivity* **3**, 317–338 (2013).
- Parlatini, V., Radua, J., Dell'Acqua, F., Leslie, A., Simmons, A., Murphy, D. G., Catani, M. & Thiebaut de Schotten, M. Functional segregation and integration within fronto-parietal networks. *NeuroImage* **146**, 367–375 (2017).
- Pascual-Leone, A., Gomez-Tortosa, E., Grafman, J., Alway, D., Nichelli, P. & Hallett, M. Induction of visual extinction by rapid-rate transcranial magnetic stimulation of parietal lobe. *Neurology* **44**, 494–494 (1994).
- Pascual-Leone, A., Walsh, V. & Rothwell, J. Transcranial magnetic stimulation in cognitive neuroscience—virtual lesion, chronometry, and functional connectivity. *Current opinion in neurobiology* **10**, 232–237 (2000).
- Patel, M., McKevitt, C., Lawrence, E., Rudd, A. & Wolfe, C. Clinical determinants of long-term quality of life after stroke. *Age and Ageing* **36**, 316–322 (2007).
- Patenaude, B., Smith, S. M., Kennedy, D. N. & Jenkinson, M. A Bayesian model of shape and appearance for subcortical brain segmentation. *Neuroimage* **56**, 907–922 (2011).

- Paz, R., Boraud, T., Natan, C., Bergman, H. & Vaadia, E. Preparatory activity in motor cortex reflects learning of local visuomotor skills. *Nature neuroscience* **6**, 882–890 (2003).
- Paz, R., Natan, C., Boraud, T., Bergman, H. & Vaadia, E. Emerging patterns of neuronal responses in supplementary and primary motor areas during sensorimotor adaptation. *Journal of Neuroscience* **25**, 10941–10951 (2005).
- Peinemann, A., Lehner, C., Conrad, B. & Siebner, H. R. Age-related decrease in paired-pulse intracortical inhibition in the human primary motor cortex. *Neuroscience letters* **313**, 33–36 (2001).
- Pekny, S. E., Izawa, J. & Shadmehr, R. Reward-dependent modulation of movement variability. *Journal of Neuroscience* **35**, 4015–4024 (2015).
- Penfield, W. & Boldrey, E. Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain: A journal of neurology* (1937).
- Petitot, P., Noonan, M. P., Bridge, H., O'Reilly, J. X. & O'Shea, J. Testing the inter-hemispheric competition account of visual extinction with combined TMS/fMRI. *Neuropsychologia* **74**, 63–73 (2015).
- Petitot, P., O'Reilly, J. X. & O'Shea, J. Towards a neuro-computational account of prism adaptation. *Neuropsychologia* (2017).
- Petrides, M. & Pandya, D. N. Projections to the frontal cortex from the posterior parietal region in the rhesus monkey. *Journal of Comparative Neurology* **228**, 105–116 (1984).
- Petroff, O. A. Book review: GABA and glutamate in the human brain. *The Neuroscientist* **8**, 562–573 (2002).
- Pisella, L., Michel, C., Grea, H., Tilikete, C., Vighetto, A. & Rossetti, Y. Preserved prism adaptation in bilateral optic ataxia: strategic versus adaptive reaction to prisms. *Experimental Brain Research* **156**, 399–408 (2004).
- Pisella, L., Rossetti, Y., Michel, C., Rode, G., Boisson, D., Pelisson, D. & Tilikete, C. Ipsidirectional impairment of prism adaptation after unilateral lesion of anterior cerebellum. *Neurology* **65**, 150–152 (2005).
- Pisella, L., Rode, G., Farne, A., Boisson, D. & Rossetti, Y. Dissociated long lasting improvements of straight-ahead pointing and line bisection tasks in two hemineglect patients. *Neuropsychologia* **40**, 327–334 (2002).
- Plow, E. B., Cattaneo, Z., Carlson, T. A., Alvarez, G. A., Pascual-Leone, A. & Battelli, L. The compensatory dynamic of inter-hemispheric interactions in visuospatial attention revealed using rTMS and fMRI. *Frontiers in human neuroscience* **8** (2014).
- Poreisz, C., Boros, K., Antal, A. & Paulus, W. Safety aspects of transcranial direct current stimulation concerning healthy subjects and patients. *Brain research bulletin* **72**, 208–214 (2007).

Porges, E. C., Woods, A. J., Lamb, D. G., Williamson, J. B., Cohen, R. A., Edden, R. A. & Harris, A. D. Impact of tissue correction strategy on GABA-edited MRS findings. *NeuroImage* **162**, 249–256 (2017).

Porges, E. C., Woods, A. J., Edden, R. A., Puts, N. A., Harris, A. D., Chen, H., Garcia, A. M., Seider, T. R., Lamb, D. G., Williamson, J. B., *et al.* Frontal gamma-aminobutyric acid concentrations are associated with cognitive performance in older adults. *Biological psychiatry: cognitive neuroscience and neuroimaging* **2**, 38–44 (2017).

Priori, A. Brain polarization in humans: a reappraisal of an old tool for prolonged non-invasive modulation of brain excitability. *Clinical Neurophysiology* **114**, 589–595 (2003).

Priori, A., Berardelli, A., Rona, S., Accornero, N. & Manfredi, M. Polarization of the human motor cortex through the scalp. *Neuroreport* **9**, 2257–2260 (1998).

Provencher, S. W. Estimation of metabolite concentrations from localized in vivo proton NMR spectra. *Magnetic resonance in medicine* **30**, 672–679 (1993).

Provencher, S. W. Automatic quantitation of localized in vivo ^1H spectra with LCModel. *NMR in Biomedicine* **14**, 260–264 (2001).

Provencher, S. LCModel & LCMgui user's manual. *Stephen Provencher Inc* (2012).

Pruim, R. H., Mennes, M., van Rooij, D., Llera, A., Buitelaar, J. K. & Beckmann, C. F. ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data. *Neuroimage* **112**, 267–277 (2015).

Purpura, D. P. & McMurtry, J. G. Intracellular activities and evoked potential changes during polarization of motor cortex. *Journal of neurophysiology* **28**, 166–185 (1965).

Radman, T., Ramos, R. L., Brumberg, J. C. & Bikson, M. Role of cortical cell type and morphology in subthreshold and suprathreshold uniform electric field stimulation in vitro. *Brain stimulation* **2**, 215–228 (2009).

Rahman, A., Lafon, B. & Bikson, M. Multilevel computational models for predicting the cellular effects of noninvasive brain stimulation. *Progress in brain research* **222**, 25–40 (2015).

Rahman, A., Lafon, B., Parra, L. C. & Bikson, M. Direct current stimulation boosts synaptic gain and cooperativity in vitro. *The Journal of Physiology* (2017).

Rahman, A., Reato, D., Arlotti, M., Gasca, F., Datta, A., Parra, L. C. & Bikson, M. Cellular effects of acute direct current stimulation: somatic and synaptic terminal effects. *The Journal of physiology* **591**, 2563–2578 (2013).

Ravizza, S. M. & Ivry, R. B. Comparison of the basal ganglia and cerebellum in shifting attention. *Journal of Cognitive Neuroscience* **13**, 285–297 (2001).

Reato, D., Rahman, A., Bikson, M. & Parra, L. C. Low-intensity electrical stimulation affects network dynamics by modulating population rate and spike timing. *Journal of Neuroscience* **30**, 15067–15079 (2010).

- Redding, G. M., Rader, S. D. & Lucas, D. R. Cognitive load and prism adaptation. *Journal of motor behavior* **24**, 238–246 (1992).
- Redding, G. M., Rossetti, Y. & Wallace, B. Applications of prism adaptation: a tutorial in theory and method. *Neuroscience & Biobehavioral Reviews* **29**, 431–444 (2005).
- Redding, G. M. & Wallace, B. Adaptive spatial alignment and strategic perceptual-motor control. *Journal of Experimental Psychology: Human Perception and Performance* **22**, 379 (1996).
- Redding, G. M. & Wallace, B. Prism exposure aftereffects and direct effects for different movement and feedback times. *Journal of Motor Behavior* **32**, 83–99 (2000).
- Redding, G. M. & Wallace, B. Calibration and alignment are separable: Evidence from prism adaptation. *Journal of motor behavior* **33**, 401–412 (2001).
- Redding, G. M. & Wallace, B. Strategic calibration and spatial alignment: a model from prism adaptation. *Journal of motor behavior* **34**, 126–138 (2002).
- Redding, G. M. & Wallace, B. Generalization of prism adaptation. *Journal of Experimental Psychology: Human Perception and Performance* **32**, 1006 (2006).
- Redding, G. M. & Wallace, B. Prism adaptation and unilateral neglect: review and analysis. *Neuropsychologia* **44**, 1–20 (2006).
- Reis, J., Schambra, H. M., Cohen, L. G., Buch, E. R., Fritsch, B., Zarahn, E., Celnik, P. A. & Krakauer, J. W. Noninvasive cortical stimulation enhances motor skill acquisition over multiple days through an effect on consolidation. *Proceedings of the National Academy of Sciences* **106**, 1590–1595 (2009).
- Reis, J., Robertson, E. M., Krakauer, J. W., Rothwell, J., Marshall, L., Gerloff, C., Wassermann, E. M., Pascual-Leone, A., Hummel, F., Celnik, P. A., *et al.* Consensus: can transcranial direct current stimulation and transcranial magnetic stimulation enhance motor learning and memory formation? *Brain stimulation* **1**, 363–369 (2008).
- Rengachary, J., d’Avossa, G., Sapir, A., Shulman, G. L. & Corbetta, M. Is the posner reaction time test more accurate than clinical tests in detecting left neglect in acute and chronic stroke? *Archives of physical medicine and rehabilitation* **90**, 2081–2088 (2009).
- Ringman, J., Saver, J., Woolson, R., Clarke, W. & Adams, H. Frequency, risk factors, anatomy, and course of unilateral neglect in an acute stroke cohort. *Neurology* **63**, 468–474 (2004).
- Rock, I., Goldberg, J. & Mack, A. Immediate correction and adaptation based on viewing a prismatically displaced scene. *Perception & Psychophysics* **1**, 351–354 (1966).
- Rode, G., Pisella, L., Petitet, P., O’Shea, J., Huchon, L., Jacquin-Courtois, S. & Rossetti, Y. Quelles stratégies de rééducation dans la négligence spatiale unilatérale? *Troubles neurocognitifs vasculaires et post-AVC: De l’évaluation à la prise en charge*, 97 (2017).

- Rode, G., Rossetti, Y. & Boisson, D. Prism adaptation improves representational neglect. *Neuropsychologia* **39**, 1250–1254 (2001).
- Rode, G., Lacour, S., Jacquin-Courtois, S., Pisella, L., Michel, C., Revol, P., Alahyane, N., Luauté, J., Gallagher, S., Halligan, P., *et al.* Long-term sensorimotor and therapeutical effects of a mild regime of prism adaptation in spatial neglect. A double-blind RCT essay. *Annals of physical and rehabilitation medicine* **58**, 40–53 (2015).
- Roller, C. A., Cohen, H. S., Kimball, K. T. & Bloomberg, J. J. Effects of normal aging on visuo-motor plasticity. *Neurobiology of aging* **23**, 117–123 (2002).
- Rorden, C. & Karnath, H.-O. A simple measure of neglect severity. *Neuropsychologia* **48**, 2758–2763 (2010).
- Rossetti, Y., Jacquin-Courtois, S., Rode, G., Ota, H., Michel, C. & Boisson, D. Does action make the link between number and space representation? Visuo-manual adaptation improves number bisection in unilateral neglect. *Psychological Science* **15**, 426–430 (2004).
- Rossetti, Y., Koga, K. & Mano, T. Prismatic displacement of vision induces transient changes in the timing of eye-hand coordination. *Attention, Perception, & Psychophysics* **54**, 355–364 (1993).
- Rossetti, Y., Rode, G., Pisella, L., Farné, A., Li, L., Boisson, D. & Perenin, M.-T. Prism adaptation to a rightward optical deviation rehabilitates left hemispatial neglect. *Nature* **395**, 166–169 (1998).
- Rossi, S., Hallett, M., Rossini, P. M. & Pascual-Leone, A. Screening questionnaire before TMS: an update. *Clinical Neurophysiology* **122**, 1686–1686 (2011).
- Rossi, S., Hallett, M., Rossini, P. M., Pascual-Leone, A., of TMS Consensus Group, S., *et al.* Safety, ethical considerations, and application guidelines for the use of transcranial magnetic stimulation in clinical practice and research. *Clinical neurophysiology* **120**, 2008–2039 (2009).
- Rossini, P. M., Burke, D., Chen, R., Cohen, L., Daskalakis, Z., Di Iorio, R., Di Lazzaro, V., Ferreri, F., Fitzgerald, P., George, M., *et al.* Non-invasive electrical and magnetic stimulation of the brain, spinal cord, roots and peripheral nerves: basic principles and procedures for routine clinical and research application. An updated report from an IFCN Committee. *Clinical Neurophysiology* **126**, 1071–1107 (2015).
- Roux, F.-E., Dufor, O., Lauwers-Cances, V., Boukhatem, L., Brauge, D., Draper, L., Lotterie, J.-A. & Démonet, J.-F. Electrostimulation mapping of spatial neglect. *Neurosurgery* **69**, 1218–1231 (2011).
- Rowland, L. M., Kontson, K., West, J., Edden, R. A., Zhu, H., Wijtenburg, S. A., Holcomb, H. H. & Barker, P. B. In vivo measurements of glutamate, GABA, and NAAG in schizophrenia. *Schizophrenia bulletin*, sbs092 (2012).
- Rozycka, A. & Liguz-Leczna, M. The space where aging acts: focus on the GABAergic synapse. *Aging Cell* (2017).
- Rushworth, M. F., Ellison, A. & Walsh, V. Complementary localization and lateralization of orienting and motor attention. *Nature neuroscience* **4**, 656–661 (2001).

- Sack, A. T., Camprodon, J. A., Pascual-Leone, A. & Goebel, R. The dynamics of interhemispheric compensatory processes in mental imagery. *Science* **308**, 702–704 (2005).
- Sack, A. T., Sperling, J. M., Prvulovic, D., Formisano, E., Goebel, R., Di Salle, F., Dierks, T. & Linden, D. E. Tracking the mind's image in the brain II: transcranial magnetic stimulation reveals parietal asymmetry in visuospatial imagery. *Neuron* **35**, 195–204 (2002).
- Sack, A., Hubl, D., Prvulovic, D., Formisano, E., Jandl, M., Zanella, F., Maurer, K., Goebel, R., Dierks, T. & Linden, D. E. J. The experimental combination of rTMS and fMRI reveals the functional relevance of parietal cortex for visuospatial functions. *Cognitive Brain Research* **13**, 85–93 (2002).
- Saevarsson, S., Kristjánsson, Á., Hildebrandt, H. & Halsband, U. Prism adaptation improves visual search in hemispatial neglect. *Neuropsychologia* **47**, 717–725 (2009).
- Sailasuta, N., Ernst, T. & Chang, L. Regional variations and the effects of age and gender on glutamate concentrations in the human brain. *Magnetic resonance imaging* **26**, 667–675 (2008).
- Sallet, J., Mars, R. B., Noonan, M. P., Neubert, F.-X., Jbabdi, S., O'Reilly, J. X., Filippini, N., Thomas, A. G. & Rushworth, M. F. The organization of dorsal frontal cortex in humans and macaques. *Journal of Neuroscience* **33**, 12255–12274 (2013).
- Samuelsson, H., Jensen, C., Ekholm, S., Naver, H. & Blomstrand, C. Anatomical and neurological correlates of acute and chronic visuospatial neglect following right hemisphere stroke. *Cortex* **33**, 271–285 (1997).
- Sandrini, M., Umiltà, C. & Rusconi, E. The use of transcranial magnetic stimulation in cognitive neuroscience: a new synthesis of methodological issues. *Neuroscience & Biobehavioral Reviews* **35**, 516–536 (2011).
- Sarmiento, C., San-Juan, D. & Prasath, V. Letter to the Editor: Brief history of transcranial direct current stimulation (tDCS): from electric fishes to microcontrollers. *Psychological medicine* **46**, 3259 (2016).
- Sarri, M., Greenwood, R., Kalra, L., Papps, B., Husain, M. & Driver, J. Prism adaptation aftereffects in stroke patients with spatial neglect: pathological effects on subjective straight ahead but not visual open-loop pointing. *Neuropsychologia* **46**, 1069–1080 (2008).
- Savignac, H. M., Corona, G., Mills, H., Chen, L., Spencer, J. P., Tzortzis, G. & Burnet, P. W. Prebiotic feeding elevates central brain derived neurotrophic factor, N-methyl-D-aspartate receptor subunits and D-serine. *Neurochemistry international* **63**, 756–764 (2013).
- Schintu, S., Patané, I., Caldano, M., Salemme, R., Reilly, K. T., Pisella, L. & Farnè, A. The asymmetrical effect of leftward and rightward prisms on intact visuospatial cognition. *Cortex* **97**, 23–31 (2017).

- Schintu, S., Pisella, L., Jacobs, S., Salemme, R., Reilly, K. T. & Farné, A. Prism adaptation in the healthy brain: the shift in line bisection judgments is long lasting and fluctuates. *Neuropsychologia* **53**, 165–170 (2014).
- Scholl, J., Kolling, N., Nelissen, N., Stagg, C. J., Harmer, C. J. & Rushworth, M. F. Excitation and inhibition in anterior cingulate predict use of past experiences. *eLife* **6**, e20365 (2017).
- Scholz, J., Tomassini, V. & Johansen-Berg, H. Individual differences in white matter microstructure in the healthy brain. *Diffusion MRI: From Quantitative Measurement to In-vivo Neuroanatomy*, 237–247 (2009).
- Serino, A., Angeli, V., Frassinetti, F. & Làdavas, E. Mechanisms underlying neglect recovery after prism adaptation. *Neuropsychologia* **44**, 1068–1078 (2006).
- Serino, A., Barbiani, M., Rinaldesi, M. L. & Làdavas, E. Effectiveness of prism adaptation in neglect rehabilitation. *Stroke* **40**, 1392–1398 (2009).
- Shadmehr, R., Huang, H. J. & Ahmed, A. A. A representation of effort in decision-making and motor control. *Current biology* **26**, 1929–1934 (2016).
- Shadmehr, R. & Mussa-Ivaldi, F. A. Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience* **14**, 3208–3224 (1994).
- Shadmehr, R., Smith, M. A. & Krakauer, J. W. Error correction, sensory prediction, and adaptation in motor control. *Annual review of neuroscience* **33**, 89–108 (2010).
- Shan, Z. Y., Wright, M. J., Thompson, P. M., McMahon, K. L., Blokland, G. G., De Zubicaray, G. I., Martin, N. G., Vinkhuyzen, A. A. & Reutens, D. C. Modeling of the hemodynamic responses in block design fMRI studies. *Journal of Cerebral Blood Flow & Metabolism* **34**, 316–324 (2014).
- Shiraishi, H., Yamakawa, Y., Itou, A., Muraki, T. & Asada, T. Long-term effects of prism adaptation on chronic neglect after stroke. *NeuroRehabilitation* **23**, 137–151 (2008).
- Shrout, P. E., Bolger, N., *et al.* Mediation in experimental and nonexperimental studies: New procedures and recommendations. *Psychological methods* **7**, 422–445 (2002).
- Shulman, G. L., Astafiev, S. V., Franke, D., Pope, D. L., Snyder, A. Z., McAvoy, M. P. & Corbetta, M. Interaction of stimulus-driven reorienting and expectation in ventral and dorsal frontoparietal and basal ganglia-cortical networks. *Journal of Neuroscience* **29**, 4392–4407 (2009).
- Shulman, G. L., Pope, D. L., Astafiev, S. V., McAvoy, M. P., Snyder, A. Z. & Corbetta, M. Right hemisphere dominance during spatial selective attention and target detection occurs outside the dorsal frontoparietal network. *Journal of Neuroscience* **30**, 3640–3651 (2010).
- Silveri, M., Misciagna, S. & Terrezza, G. Right side neglect in right cerebellar lesion. *Journal of Neurology, Neurosurgery & Psychiatry* **71**, 114–117 (2001).

- Simmons, G. & Demiris, Y. Optimal robot arm control using the minimum variance model. *Journal of Field Robotics* **22**, 677–690 (2005).
- Smith, M. A., Ghazizadeh, A. & Shadmehr, R. Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* **4**, e179 (2006).
- Smith, M., Keel, J., Greenberg, B., Adams, L., Schmidt, P., Rubinow, D. & Wassermann, E. M. Menstrual cycle effects on cortical excitability. *Neurology* **53**, 2069–2069 (1999).
- Smith, S. M. Fast robust automated brain extraction. *Human brain mapping* **17**, 143–155 (2002).
- Smith, S. M., Nichols, T. E., Vidaurre, D., Winkler, A. M., Behrens, T. E., Glasser, M. F., Ugurbil, K., Barch, D. M., Van Essen, D. C. & Miller, K. L. A positive-negative mode of population covariation links brain connectivity, demographics and behavior. *Nature neuroscience* **18**, 1565–1567 (2015).
- Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E., Johansen-Berg, H., Bannister, P. R., De Luca, M., Drobnjak, I., Flitney, D. E., *et al.* Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* **23**, S208–S219 (2004).
- Smith, S. M., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T. E., Mackay, C. E., Watkins, K. E., Ciccarelli, O., Cader, M. Z., Matthews, P. M., *et al.* Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *Neuroimage* **31**, 1487–1505 (2006).
- Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R., *et al.* Correspondence of the brain’s functional architecture during activation and rest. *Proceedings of the National Academy of Sciences* **106**, 13040–13045 (2009).
- Soares, J. M., Magalhães, R., Moreira, P. S., Sousa, A., Ganz, E., Sampaio, A., Alves, V., Marques, P. & Sousa, N. A hitchhiker’s guide to functional magnetic resonance imaging. *Frontiers in neuroscience* **10** (2016).
- Sobel, M. E. Asymptotic confidence intervals for indirect effects in structural equation models. *Sociological methodology* **13**, 290–312 (1982).
- Sober, S. J. & Sabes, P. N. Flexible strategies for sensory integration during motor planning. *Nature neuroscience* **8**, 490–497 (2005).
- Song, W., Du, B., Xu, Q., Hu, J., Wang, M. & Luo, Y. Low-frequency transcranial magnetic stimulation for visual spatial neglect: a pilot study. *Journal of Rehabilitation Medicine* **41**, 162–165 (2009).
- Sparing, R., Thimm, M., Hesse, M., Küst, J., Karbe, H. & Fink, G. Bidirectional alterations of interhemispheric parietal balance by non-invasive cortical stimulation. *Brain* **132**, 3011–3020 (2009).
- Stagg, C. J., Bachtiar, V. & Johansen-Berg, H. The role of GABA in human motor learning. *Current Biology* **21**, 480–484 (2011).

- Stagg, C. J., Bachtiar, V. & Johansen-Berg, H. What are we measuring with GABA magnetic resonance spectroscopy? *Communicative & integrative biology* **4**, 573–575 (2011).
- Stagg, C. J., Best, J. G., Stephenson, M. C., O’Shea, J., Wylezinska, M., Kincses, Z. T., Morris, P. G., Matthews, P. M. & Johansen-Berg, H. Polarity-sensitive modulation of cortical neurotransmitters by transcranial stimulation. *Journal of Neuroscience* **29**, 5202–5206 (2009).
- Stagg, C. J. & Nitsche, M. A. Physiological basis of transcranial direct current stimulation. *The Neuroscientist* **17**, 37–53 (2011).
- Stagg, C. J., Bachtiar, V., Amadi, U., Gudberg, C. A., Ilie, A. S., Sampaio-Baptista, C., O’Shea, J., Woolrich, M., Smith, S. M., Filippini, N., *et al.* Local GABA concentration is related to network-level resting functional connectivity. *Elife* **3**, e01465 (2014).
- Stagg, C. & Rothman, D. L. *Magnetic resonance spectroscopy: tools for neuroscience research and emerging clinical applications* (Academic Press, 2013).
- Stagg, C., Bestmann, S., Constantinescu, A., Moreno Moreno, L., Allman, C., Mekle, R., Woolrich, M., Near, J., Johansen-Berg, H. & Rothwell, J. Relationship between physiological measures of excitability and levels of glutamate and GABA in the human motor cortex. *The Journal of physiology* **589**, 5845–5855 (2011).
- Stagg, C., Jayaram, G., Pastor, D., Kincses, Z., Matthews, P. & Johansen-Berg, H. Polarity and timing-dependent effects of transcranial direct current stimulation in explicit motor learning. *Neuropsychologia* **49**, 800–804 (2011).
- Steinberg, H. Letter to the editor: transcranial direct current stimulation (tDCS) has a history reaching back to the 19th century. *Psychological medicine* **43**, 669 (2013).
- Stern, Y., Mayeux, R., Hermann, A. & Rosen, J. Prism adaptation in Parkinson’s disease. *Journal of Neurology, Neurosurgery & Psychiatry* **51**, 1584–1587 (1988).
- Stewart, L., Walsh, V. & Rothwell, J. Motor and phosphene thresholds: a transcranial magnetic stimulation correlation study. *Neuropsychologia* **39**, 415–419 (2001).
- Stokes, M. G., Barker, A. T., Dervinis, M., Verbruggen, F., Maizey, L., Adams, R. C. & Chambers, C. D. Biophysical determinants of transcranial magnetic stimulation: effects of excitability and depth of targeted area. *Journal of neurophysiology* **109**, 437–444 (2013).
- Stone, J. M., Dietrich, C., Edden, R., Mehta, M. A., De Simoni, S., Reed, L. J., Krystal, J. H., Nutt, D. & Barker, G. J. Ketamine effects on brain GABA and glutamate levels with 1H-MRS: relationship to ketamine-induced psychopathology. *Molecular psychiatry* **17** (2012).
- Stone, S., Halligan, P. & Greenwood, R. The incidence of neglect phenomena and related disorders in patients with an acute right or left hemisphere stroke. *Age and Ageing* **22**, 46–52 (1993).
- Stoodley, C. J. & Schmahmann, J. D. Functional topography in the human cerebellum: a meta-analysis of neuroimaging studies. *Neuroimage* **44**, 489–501 (2009).

- Strata, P. The emotional cerebellum. *The Cerebellum* **14**, 570–577 (2015).
- Striemer, C. L. & Danckert, J. A. Through a prism darkly: re-evaluating prisms and neglect. *Trends in cognitive sciences* **14**, 308–316 (2010).
- Striemer, C. L., Russell, K. & Nath, P. Prism adaptation magnitude has differential influences on perceptual versus manual responses. *Experimental brain research* **234**, 2761–2772 (2016).
- Sumitani, M., Rossetti, Y., Shibata, M., Matsuda, Y., Sakaue, G., Inoue, T., Mashimo, T. & Miyauchi, S. Prism adaptation to optical deviation alleviates pathologic pain. *Neurology* **68**, 128–133 (2007).
- Sunwoo, H., Kim, Y.-H., Chang, W. H., Noh, S., Kim, E.-J. & Ko, M.-H. Effects of dual transcranial direct current stimulation on post-stroke unilateral visuospatial neglect. *Neuroscience letters* **554**, 94–98 (2013).
- Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction* **1** (MIT press Cambridge, 1998).
- Szczepanski, S. M. & Kastner, S. Shifting attentional priorities: control of spatial attention through hemispheric competition. *Journal of Neuroscience* **33**, 5411–5421 (2013).
- Szczepanski, S. M., Konen, C. S. & Kastner, S. Mechanisms of spatial attention control in frontal and parietal cortex. *Journal of Neuroscience* **30**, 148–160 (2010).
- Takahashi, M., Hackney, D. B., Zhang, G., Wehrli, S. L., Wright, A. C., O'Brien, W. T., Uematsu, H., Wehrli, F. W. & Selzer, M. E. Magnetic resonance microimaging of intraaxonal water diffusion in live excised lamprey spinal cord. *Proceedings of the National Academy of Sciences* **99**, 16192–16196 (2002).
- Tanaka, H., Krakauer, J. W. & Sejnowski, T. J. Generalization and multirate models of motor adaptation. *Neural computation* **24**, 939–966 (2012).
- Tanaka, H., Sejnowski, T. J. & Krakauer, J. W. Adaptation to visuomotor rotation through interaction between posterior parietal and motor cortical areas. *Journal of neurophysiology* **102**, 2921–2932 (2009).
- Tanaka, J.-i., Horiike, Y., Matsuzaki, M., Miyazaki, T., Ellis-Davies, G. C. & Kasai, H. Protein synthesis and neurotrophin-dependent structural plasticity of single dendritic spines. *Science* **319**, 1683–1687 (2008).
- Tavor, I., Jones, O. P., Mars, R., Smith, S., Behrens, T. & Jbabdi, S. Task-free MRI predicts individual differences in brain activity during task performance. *Science* **352**, 216–220 (2016).
- Taylor, J. A., Krakauer, J. W. & Ivry, R. B. Explicit and implicit contributions to learning in a sensorimotor adaptation task. *Journal of Neuroscience* **34**, 3023–3032 (2014).
- Templeton, W. B., Howard, I. P. & Wilkinson, D. A. Additivity of components of prismatic adaptation. *Attention, Perception, & Psychophysics* **15**, 249–257 (1974).

Ten Brink, A. F., Visser-Meily, J. M., Schut, M. J., Kouwenhoven, M., Eijssackers, A. L. & Nijboer, T. C. Prism Adaptation in Rehabilitation? No Additional Effects of Prism Adaptation on Neglect Recovery in the Subacute Phase Poststroke: A Randomized Controlled Trial. *Neurorehabilitation and neural repair*, 1545968317744277 (2017).

Theodore, W., Hunter, K., Chen, R., Vega-Bermudez, F., Boroojerdi, B., Reeves-Tyer, P., Werhahn, K., Kelley, K. & Cohen, L. Transcranial magnetic stimulation for the treatment of seizures A controlled study. *Neurology* **59**, 560–562 (2002).

Thiebaut De Schotten, M., Dell’Acqua, F., Forkel, S. J., Simmons, A., Vergani, F., Murphy, D. G. & Catani, M. A lateralized brain network for visuospatial attention. *Nature neuroscience* **14**, 1245–1246 (2011).

Thiebaut de Schotten, M., Tomaiuolo, F., Aiello, M., Merola, S., Silvetti, M., Lecce, F., Bartolomeo, P. & Doricchi, F. Damage to white matter pathways in subacute and chronic spatial neglect: a group study and 2 single-case studies with complete virtual “in vivo” tractography dissection. *Cerebral Cortex* **24**, 691–706 (2012).

Thiebaut de Schotten, M., Urbanski, M., Duffau, H., Volle, E., Lévy, R., Dubois, B. & Bartolomeo, P. Direct evidence for a parietal-frontal pathway subserving spatial awareness in humans. *Science* **309**, 2226–2228 (2005).

Thiebaut de Schotten, M., Bizzi, A., Dell’Acqua, F., Allin, M., Walshe, M., Murray, R., Williams, S. C., Murphy, D. G., Catani, M., *et al.* Atlasing location, asymmetry and inter-subject variability of white matter tracts in the human brain with MR diffusion tractography. *Neuroimage* **54**, 49–59 (2011).

Thiebaut de Schotten, M., Dell’Acqua, F., Ratiu, P., Leslie, A., Howells, H., Cabanis, E., Iba-Zizen, M., Plaisant, O., Simmons, A., Dronkers, N., *et al.* From Phineas Gage and Monsieur Leborgne to HM: revisiting disconnection syndromes. *Cerebral Cortex* **25**, 4812–4827 (2015).

Thielscher, A., Antunes, A. & Saturnino, G. B. *Field modeling for transcranial magnetic stimulation: A useful tool to understand the physiological effects of TMS? Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE* (2015), 222–225.

Tisell, A., Leinhard, O. D., Warntjes, J. B. M. & Lundberg, P. Procedure for quantitative ¹H magnetic resonance spectroscopy and tissue characterization of human brain tissue based on the use of quantitative magnetic resonance imaging. *Magnetic resonance in medicine* **70**, 905–915 (2013).

Tosoni, A., Shulman, G. L., Pope, A. L., McAvoy, M. P. & Corbetta, M. Distinct representations for shifts of spatial attention and changes of reward contingencies in the human brain. *Cortex* **49**, 1733–1749 (2013).

Tremblay, S., Beaulé, V., Lepage, J.-F. & Théoret, H. Anodal transcranial direct current stimulation modulates GABAB-related intracortical inhibition in the M1 of healthy individuals. *Neuroreport* **24**, 46–50 (2013).

- Trepel, C. & Racine, R. J. GABAergic modulation of neocortical long-term potentiation in the freely moving rat. *Synapse* **35**, 120–128 (2000).
- Treue, S. Attentional modulation of visual motion processing in cortical areas MT and MST (1996).
- Treue, S. & Maunsell, J. H. Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. *Journal of Neuroscience* **19**, 7591–7602 (1999).
- Trommershäuser, J., Maloney, L. T. & Landy, M. S. Statistical decision theory and the selection of rapid, goal-directed movements. *JOSA A* **20**, 1419–1433 (2003).
- Tu, Y.-K. & Gilthorpe, M. S. Revisiting the relation between change and initial value: a review and evaluation. *Statistics in medicine* **26**, 443–457 (2007).
- Turton, A. J., O’Leary, K., Gabb, J., Woodward, R. & Gilchrist, I. D. A single blinded randomised controlled pilot trial of prism adaptation for improving self-care in stroke patients with neglect. *Neuropsychological Rehabilitation* **20**, 180–196 (2010).
- Umarova, R. M., Saur, D., Kaller, C. P., Vry, M.-S., Glauche, V., Mader, I., Hennig, J. & Weiller, C. Acute visual neglect and extinction: distinct functional state of the visuospatial attention system. *Brain* **134**, 3310–3325 (2011).
- Umarova, R. M., Saur, D., Schnell, S., Kaller, C. P., Vry, M.-S., Glauche, V., Rijntjes, M., Hennig, J., Kiselev, V. & Weiller, C. Structural connectivity for visuospatial attention: significance of ventral pathways. *Cerebral cortex* **20**, 121–129 (2009).
- Underwood, E. Cadaver study challenges brain stimulation methods. *Science* **352**, 397–397 (2016).
- Urbanski, M., De Schotten, M. T., Rodrigo, S., Oppenheim, C., Touzé, E., Méder, J.-F., Moreau, K., Loeper-Jeny, C., Dubois, B. & Bartolomeo, P. DTI-MR tractography of white matter damage in stroke patients with neglect. *Experimental brain research* **208**, 491–505 (2011).
- Valero-Cabré, A., Amengual, J., Stengel, C., Pascual-Leone, A. & Coubard, O. A. Transcranial magnetic stimulation in basic and clinical neuroscience: a comprehensive review of fundamental principles and novel insights. *Neuroscience & Biobehavioral Reviews* (2017).
- Valero-Cabré, A., Payne, B. R. & Pascual-Leone, A. Opposite impact on 14C-2-deoxyglucose brain metabolism following patterns of high and low frequency repetitive transcranial magnetic stimulation in the posterior parietal cortex. *Experimental brain research* **176**, 603–615 (2007).
- Vallar, G., Bello, L., Bricolo, E., Castellano, A., Casarotti, A., Falini, A., Riva, M., Fava, E. & Papagno, C. Cerebral correlates of visuospatial neglect: a direct cerebral stimulation study. *Human brain mapping* **35**, 1334–1350 (2014).
- Vallar, G. & Perani, D. The anatomy of unilateral neglect after right-hemisphere stroke lesions. A clinical/CT-scan correlation study in man. *Neuropsychologia* **24**, 609–622 (1986).

- Van Den Heuvel, M. P., Pol, H. & Hilleke, E. Specific somatotopic organization of functional connections of the primary motor network during resting state. *Human brain mapping* **31**, 631–644 (2010).
- Van Beers, R. J., Sittig, A. C. & Gon, J. J. D. v. d. Integration of proprioceptive and visual position-information: An experimentally supported model. *Journal of neurophysiology* **81**, 1355–1364 (1999).
- Van Beers, R. J., Sittig, A. C. & van der Gon Denier, J. J. How humans combine simultaneous proprioceptive and visual position information. *Experimental Brain Research* **111**, 253–261 (1996).
- Vangkilde, S. & Habekost, T. Finding Wally: prism adaptation improves visual search in chronic neglect. *Neuropsychologia* **48**, 1994–2004 (2010).
- Vaswani, P. A. & Shadmehr, R. Decay of motor memories in the absence of error. *Journal of neuroscience* **33**, 7700–7709 (2013).
- Verdon, V., Schwartz, S., Lovblad, K.-O., Hauert, C.-A. & Vuilleumier, P. Neuroanatomy of hemispatial neglect and its functional components: a study using voxel-based lesion-symptom mapping. *Brain* **133**, 880–894 (2009).
- Verhagen, M., Van Der Meij, A., Van Deurzen, P., Janzing, J., Arias-Vasquez, A., Buitelaar, J. & Franke, B. Meta-analysis of the BDNF Val66Met polymorphism in major depressive disorder: effects of gender and ethnicity. *Molecular psychiatry* **15**, 260–271 (2010).
- Vetter, P. & Wolpert, D. M. Context estimation for sensorimotor control. *Journal of Neurophysiology* **84**, 1026–1034 (2000).
- Von Helmholtz, H. Treatise on Physiological Optics Vol. III (1867).
- Vuilleumier, P., Sagiv, N., Hazeltine, E., Poldrack, R. A., Swick, D., Rafal, R. D. & Gabrieli, J. D. Neural fate of seen and unseen faces in visuospatial neglect: a combined event-related functional MRI and event-related potential study. *Proceedings of the National Academy of Sciences* **98**, 3495–3500 (2001).
- Vuilleumier, P., Schwartz, S., Verdon, V., Maravita, A., Hutton, C., Husain, M. & Driver, J. Abnormal attentional modulation of retinotopic cortex in parietal patients with spatial neglect. *Current Biology* **18**, 1525–1529 (2008).
- Wagner, S., Rampersad, S., Aydin, ü., Vorwerk, J., Oostendorp, T., Neuling, T., Herrmann, C., Stegeman, D. & Wolters, C. Investigation of tDCS volume conduction effects in a highly realistic head model. *Journal of neural engineering* **11**, 016002 (2013).
- Warach, S., Chien, D., Li, W., Ronthal, M. & Edelman, R. Fast magnetic resonance diffusion-weighted imaging of acute human stroke. *Neurology* **42**, 1717–1717 (1992).
- Wassermann, E. M. Risk and safety of repetitive transcranial magnetic stimulation: report and suggested guidelines from the International Workshop on the Safety of Repetitive Transcranial Magnetic Stimulation, June 5–7, 1996. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section* **108**, 1–16 (1998).

- Wassermann, E., Epstein, C. & Ziemann, U. *Oxford handbook of transcranial stimulation* (Oxford University Press, 2008).
- Watanabe, S. & Amimoto, K. Generalization of prism adaptation for wheelchair driving task in patients with unilateral spatial neglect. *Archives of physical medicine and rehabilitation* **91**, 443–447 (2010).
- Wei, K. & Körding, K. Uncertainty of feedback and state estimation determines the speed of motor adaptation. *Frontiers in computational neuroscience* **4** (2010).
- Weinberg, J., Diller, L., Gordon, W., Gerstman, L., Lieberman, A., Lakin, P., Hodges, G. & Ezrachi, O. Training sensory awareness and spatial organization in people with right brain damage. *Archives of physical medicine and rehabilitation* **60**, 491–496 (1979).
- Weinberg, J., Diller, L., Gordon, W., Gerstman, L., Lieberman, A., Lakin, P., Hodges, G. & Ezrachi, O. Visual scanning training effect on reading-related tasks in acquired right brain damage. *Archives of physical medicine and rehabilitation* **58**, 479–486 (1977).
- Weiner, M. J., Hallett, M. & Funkenstein, H. H. Adaptation to lateral displacement of vision in patients with lesions of the central nervous system. *Neurology* **33**, 766–766 (1983).
- Weintraub, S. & Mesulam, M.-M. Right cerebral dominance in spatial attention: Further evidence based on ipsilateral neglect. *Archives of neurology* **44**, 621–625 (1987).
- Welch, R. B. & Goldstein, G. Prism adaptation and brain damage. *Neuropsychologia* **10**, 387–394 (1972).
- Welch, R. B., Widawski, M. H., Harrington, J. & Warren, D. H. An examination of the relationship between visual capture and prism adaptation. *Attention, Perception, & Psychophysics* **25**, 126–132 (1979).
- Whittall, K. P., Mackay, A. L., Graeb, D. A., Nugent, R. A., Li, D. K. & Paty, D. W. In vivo measurement of T2 distributions and water contents in normal human brain. *Magnetic resonance in medicine* **37**, 34–43 (1997).
- Wigström, H. & Gustafsson, B. Facilitated induction of hippocampal long-lasting potentiation during blockade of inhibition. *Nature* **301**, 603–604 (1983).
- Wilcoxon, F. Individual comparisons by ranking methods. *Biometrics bulletin* **1**, 80–83 (1945).
- Wilson, B., Cockburn, J. & Halligan, P. Development of a behavioral test of visuospatial neglect. *Archives of physical medicine and rehabilitation* **68**, 98–102 (1987).
- Winkler, A. M., Ridgway, G. R., Douaud, G., Nichols, T. E. & Smith, S. M. Faster permutation inference in brain imaging. *Neuroimage* **141**, 502–516 (2016).
- Winkler, A. M., Ridgway, G. R., Webster, M. A., Smith, S. M. & Nichols, T. E. Permutation inference for the general linear model. *Neuroimage* **92**, 381–397 (2014).

- Winkler, A. M., Webster, M. A., Brooks, J. C., Tracey, I., Smith, S. M. & Nichols, T. E. Non-parametric combination and related permutation tests for neuroimaging. *Human brain mapping* **37**, 1486–1511 (2016).
- Wise, S., Moody, S., Blomstrom, K. & Mitz, A. Changes in motor cortical activity during visuomotor adaptation. *Experimental Brain Research* **121**, 285–299 (1998).
- Wojciulik, E., Rorden, C., Clarke, K., Husain, M. & Driver, J. Group study of an “undercover” test for visuospatial neglect: invisible cancellation can reveal more neglect than standard cancellation. *Journal of Neurology, Neurosurgery & Psychiatry* **75**, 1356–1358 (2004).
- Wolf, P. A., D’agostino, R. B., O’neal, M. A., Sytkowski, P., Kase, C. S., Belanger, A. J. & Kannel, W. B. Secular trends in stroke incidence and mortality. The Framingham Study. *Stroke* **23**, 1551–1555 (1992).
- Wolfe, C. D. The impact of stroke. *British medical bulletin* **56**, 275–286 (2000).
- Wolpe, N., Ingram, J. N., Tsvetanov, K. A., Geerligs, L., Kievit, R. A., Henson, R. N., Wolpert, D. M., Tyler, L. K., Brayne, C., Bullmore, E., *et al.* Ageing increases reliance on sensorimotor prediction through structural and functional differences in frontostriatal circuits. *Nature communications* **7**, ncomms13034 (2016).
- Wolpert, D. M., Diedrichsen, J. & Flanagan, J. R. Principles of sensorimotor learning. *Nature Reviews Neuroscience* **12**, 739–751 (2011).
- Wolpert, D. M. & Kawato, M. Multiple paired forward and inverse models for motor control. *Neural networks* **11**, 1317–1329 (1998).
- Woods, A. J., Antal, A., Bikson, M., Boggio, P. S., Brunoni, A. R., Celnik, P., Cohen, L. G., Fregni, F., Herrmann, C. S., Kappenman, E. S., *et al.* A technical guide to tDCS, and related non-invasive brain stimulation tools. *Clinical Neurophysiology* **127**, 1031–1048 (2016).
- Woolrich, M. W., Jbabdi, S., Patenaude, B., Chappell, M., Makni, S., Behrens, T., Beckmann, C., Jenkinson, M. & Smith, S. M. Bayesian analysis of neuroimaging data in FSL. *Neuroimage* **45**, S173–S186 (2009).
- Woolrich, M. W., Ripley, B. D., Brady, M. & Smith, S. M. Temporal autocorrelation in univariate linear modeling of FMRI data. *Neuroimage* **14**, 1370–1386 (2001).
- Yamaguchi, S., Tsuchiya, H. & Kobayashi, S. Visuospatial attention shift and motor responses in cerebellar disorders. *Journal of Cognitive Neuroscience* **10**, 95–107 (1998).
- Yamamoto, M. & Ando, H. Probabilistic models of state estimation predict visuo-motor transformations during prism adaptation. *Visual neuroscience* **29**, 119–129 (2012).
- Yang, W., Liu, T.-T., Song, X.-B., Zhang, Y., Li, Z.-H., Cui, Z.-H., Hao, Q., Liu, H. L., Lei, C. L. & Liu, J. Comparison of different stimulation parameters of repetitive transcranial magnetic stimulation for unilateral spatial neglect in stroke patients. *Journal of the neurological sciences* **359**, 219–225 (2015).

- Yavari, F., Towhidkhah, F. & Ahmadi-Pajouh, M. A. Are fast/slow process in motor adaptation and forward/inverse internal model two sides of the same coin? *Medical hypotheses* **81**, 592–600 (2013).
- Yi, Y. G., Chun, M. H., Do, K. H., Sung, E. J., Kwon, Y. G. & Kim, D. Y. The effect of transcranial direct current stimulation on neglect syndrome in stroke patients. *Annals of rehabilitation medicine* **40**, 223–229 (2016).
- Yoshiike, Y., Kimura, T., Yamashita, S., Furudate, H., Mizoroki, T., Murayama, M. & Takashima, A. GABAA Receptor-mediated acceleration of aging-associated memory decline in APP/PS1 mice and its pharmacological treatment by picrotoxin. *PloS one* **3**, e3029 (2008).
- Yousry, T., Schmid, U., Alkadhi, H., Schmidt, D., Peraud, A., Buettner, A. & Winkler, P. Localization of the motor hand area to a knob on the precentral gyrus. A new landmark. *Brain: a journal of neurology* **120**, 141–157 (1997).
- Zarahn, E., Aguirre, G. K. & D’Esposito, M. Empirical analyses of BOLD fMRI statistics. *Neuroimage* **5**, 179–197 (1997).
- Zarahn, E., Weston, G. D., Liang, J., Mazzoni, P. & Krakauer, J. W. Explaining savings for visuomotor adaptation: linear time-invariant state-space models are not sufficient. *Journal of neurophysiology* **100**, 2537–2548 (2008).
- Zhao, H., Qiao, L., Fan, D., Zhang, S., Turel, O., Li, Y., Li, J., Xue, G., Chen, A. & He, Q. Modulation of Brain Activity with Noninvasive Transcranial Direct Current Stimulation (tDCS): Clinical Applications and Safety Concerns. *Frontiers in Psychology* **8** (2017).