Examining functional plasticity in the human brain



Adam Steel

A thesis submitted for the degree of Doctor of Philosophy

Lincoln College

University of Oxford

October 2018

Table of contents

1 Introduction 1

1.1 Skill: a curious thing 1

1.1.1 What defines skilled behaviour? 3

1.1.2 How to measure skill learning? 4

1.1.3 Motor skill learning 5

1.2 The motor system 10

1.2.1 Primary motor cortex 11

1.2.2 Premotor cortex 13

1.2.3 Motor areas on the medial wall 17

1.2.4 The basal ganglia 20

1.2.5 The cerebellum 23

1.2.6 Summary 25

1.3 Neural mechanisms for learning 25

1.3.1 Activity dependent plasticity 25

1.3.2 Synaptic modification for learning 26

1.3.3 Population tuning for learning 27

1.3.4 The bigger picture 33

1.3.5 Network-level activity changes 36

1.3.6 Summary 38

1.4 Functional plasticity 38

1.4.1 Functional plasticity after training 40

1.5 Research questions 41

1.6 Aims 42

2 Methods 44

2.1 Tasks for studying sequence learning 44

2.1.1 The serial reaction time task 45

2.1.2 The force tracking task 46

2.1.3 Comparison across tasks 48

2.2 Measuring neural function in humans 48

2.2.1 Concepts in magnetic resonance 49

2.2.2 Magnetic resonance for imaging 51

2.2.3 Functional magnetic resonance imaging 53

2.2.4 FMRI preprocessing: A tale of two pipelines 61

2.2.5 Resting-state group analysis ... *66*
2.3 Magnetic resonance spectroscopic imaging ... 67
2.3.1 Magnetic resonance spectroscopy: the basics ... *69*
2.3.2 Magnetic resonance spectroscopy: localization ... *70*
2.3.3 Magnetic resonance spectroscopic imaging ... *71*
2.3.4 K-space sampling ... *71*
2.3.5 Metabolite cycling for noise correction ... *74*
2.3.6 Further preprocessing of MRSI ... *75*
2.3.7 Metabolite quantification ... *76*
3 Evidence for task-specific functional plasticity in humans ... 79
3.1 Introduction ... 79
3.2 Methods ... 81
3.2.1 Participants ... *81*
3.2.2 Study design ... *82*
3.2.3 Behavioral tasks ... *83*
3.2.4 Behavioural data analysis ... *86*
3.2.5 MRI acquisition ... *88*
3.2.6 FMRI data pre-processing ... *90*
3.2.7 ROI localization ... *92*
3.2.8 Resting-state functional connectivity analyses ... *93*
3.3 Results ... 95
3.3.1 Impact of training on behavioural performance ... *96*
3.3.2 Task-specific effects of training ... *97*
3.3.3 LOSO analysis to determine effect magnitude ... *101*
3.3.4 Relationship between regions showing functional connectivity change and behaviour ... *104*
3.3.5 Functional connectivity changes reflecting behavioural performance ... *105*
3.3.6 Non-specific functional connectivity changes ... *106*
3.4 Discussion ... 109
3.4.1 Task-specific modulation after visuo-spatial training ... *109*
3.4.2 Sensorimotor cortex connectivity changes after motor-sequence learning ... *110*
3.4.3 Non-specific effects leading to functional connectivity change ... *112*
3.4.4 Limitations ... *114*
3.5 Conclusions ... 115
4 The impact of reward and punishment depends on task demands ... 117
4.1 Introduction ... 117
4.2 Methods ... 120
4.2.1 Participants ... *123*
4.2.2 Serial reaction time task ... *124*
4.2.3 Force tracking task ... *125*
4.2.4 Feedback ... *126*
4.2.5 Statistical analysis ... *129*
4.3 Results ... 129

4.3.1 Punishment improves online series reaction time task performance 129
4.3.2 Punishment impairs force tracking task performance 137
4.3.3 Conscious recall on the serial reaction time task 139
4.3.4 Earnings 140
4.4 Discussion 141
4.4.1 Limitations 144
4.5 Conclusions 145
5 Intentionality modulates the impact of feedback on performance 146
5.1 Introduction 146
5.2 Methods 148
5.2.1 Participants 151
5.2.2 Serial reaction time task 151
5.2.3 Intentionality manipulation 152
5.2.4 Feedback 153
5.2.5 Conscious recall 154
5.2.6 Statistical analysis 156
5.3 Results 157
5.3.1 Experiment 1 – Unintentional learning 157
5.3.2 Experiment 2 – Intentional learning 163
5.3.3 Punishment specifically benefits unintentional learning 166
5.4 Discussion 168
5.4.1 Limitations 173
5.5 Conclusions 174
6 Differential impact of reward and punishment on functional connectivity after skill learning 175
6.1 Introduction 175
6.2 Methods 177
6.2.1 Participants 178
6.2.2 MRI acquisition 179
6.2.3 Resting state fMRI preprocessing 180
6.2.4 Premotor cortex functional connectivity analysis 181
6.3 Results 182
6.3.1 Reward and punishment evoked dissociable changes in PMC connectivity after SRTT training 182
6.3.2 Punishment promotes PMC-striatum connectivity after training on the FTT 184
6.3.3 Influence of reward and punishment on PMC connectivity after training depends on task 187
6.4 Discussion 189
6.4.1 Differential effects of reward and punishment on PMC connectivity after SRTT training

190
6.4.2 Punishment influences PMC functional connectivity after FTT training 192
6.4.3 The impact of feedback on functional connectivity depends on task 192
6.4.4 Limitations 195
6.5 Conclusions 196
7 Towards a link between network-level neurochemical concentration, functional connectivity and sensorimotor behaviour 198
7.1 Introduction 199
7.2 Methods 201
7.2.1 Procedure 201
7.2.2 Participants 202
7.3 MRI acquisition 202
7.3.1 T1-images 202
7.3.2 Magnetic resonance spectroscopic images 203
7.3.3 MRSI alignment to T1 image 206
7.3.4 Resting-state fMRI 208
7.3.5 Behavioural tasks 209
7.3.6 MRI analysis 212
7.4 Results 216
7.4.1 DW-CRT enables highly-reproducible GABA measurements at 3-Tesla 216
7.4.2 MRSI reliability 221
7.4.3 Characterising relationships between neurochemical concentrations and behaviour 223
7.5 Discussion 231
7.5.2 Limitations 236
7.5.3 Open questions and future directions 236
7.6 Conclusions 237
8 General discussion 238
8.1 Summary of findings 238
8.1.1 Overall summary 240
8.2 General limitations 241
8.2.1 Limitations of resting-state functional connectivity 241
8.2.2 Limitations of behavioural tasks 245
8.3 Future directions and open questions 246
8.3.1 How does neurochemical distribution across multiple regions predict task performance? 246
8.3.2 How does learning impact neurochemical concentration across brain? 247
8.3.3 Strategies employed when learning with reward and punishment 248
8.3.4 Naturalistic investigations in motor learning 249

Acronyms

AFNI	Analysis of Functional NeuroImages
AM	Morning
ANATICOR	Anatomy-based correction (noise correction)
ANOVA	Analysis of variance
AST	Action selection task
AUC	Area under the curve
B0	Static magnetic field
B1	Oscillating magnetic field
BMI	Brain machine interface
BOLD	Blood oxygen level dependent (signal)
C	Carbon
CBV	Cerebral blood volume
Cho	Choline
Cr	Creatine
CRLB	Cramer-Rao Lower Bound
CRT	Concentric ring trajectory
CSF	Cerebrospinal fluid
DW	Density weighted
EPI	Echo planar imaging
EPSI	Echo planar spectroscopic imaging
FDR	False discovery rate
FID	Free induction decay
FIX-ICA	FSL X-noisifier (noise correction)
fMRI	Functional magnetic resonance imaging
FSL	FMRIB Software Library
FTT	Force tracking task
FWHM	Full-width half-maximum
GABA	gamma-Aminobutyric acid
GCOR	Global correlation
Glu	Glutamate
Glx	Glutamate+glutamine
H	Hydrogen
HbT	Blood oxygenation
HLSVD	Hankel-Lanczos singular value decomposition

ICA	Independent components analysis
IPS	Intraparietal sulcus
LME	Linear mixed effects model
LOSO	Leave-one-subject-out
LTD	Long-term depression
LTP	Long-term potentiation
LW	Linewidth
M0	Net magnetisation
me-ICA	Multi-echo independent components analysis (noise correction)
MEGA-PRESS	Mescher-Garwood Point resolved spectroscopy
MR	Magnetic resonance
MRI	Magnetic resonance imaging
MRSI	Magnetic resonance spectroscopic imaging
MTL	Medial temporal lobe
NAA	N-acetylaspartate
NIH	National Institutes of Health
NUFFT	Non-uniform fast Fourier transform
OVS	Outer volume suppression
P	Phosphorous
PCA	Principle components analysis
PM	Afternoon
PMC	Premotor cortex
PRESS	Point resolved spectroscopy
R	Random-sequence blocks (Section 5.2.2)
rCBF	Regional cerebral blood flow
RE	Radially equidistant
RETROICOR	Retrospective correction (physiological noise correction)
RF	Radio frequency
ROI	Region of interest
RT	Response time
RVT	Respiration volume over time
S	Fixed-sequence blocks (Section 5.2.2)
SAR	Specific absorption rate
semi-LASER	Semi-localisation by adiabatic selective refocusing
SMA	Supplementary motor area
SPECIAL	Spin echo full-intensity acquired localised
SRF	Spatial response function
SRTT	Serial reaction time task
STEAM	Stimulated echo acquisition mode
svMRS	Single voxel magnetic resonance spectroscopy
T1	Magnetisation signal parallel to B0
T2	Magnetisation signal in the transverse plane
tCho	Choline+Phosphocholine

TE	Echo time
TMS	Transcranial magnetic stimulation
tNAA	N-acetyl aspartate+N-acetylasparglutamate
TOJT	Temporal order judgement task
TR	Repetition time
VTA	Ventral tegmental area
YKIWYSI	You know it when you see it

List of figures

Figure 1.1. The Marina Bay Street Circuit, looking south toward the Marina Bay Sands hotel. 1

Figure 1.2. Classical archers exhibit a highly specific skill. 3

Figure 1.3. Perceptual skill learning assessed with a texture discrimination task. 5

Figure 1.4. The speed-accuracy tradeoff demonstrated by Woodworth, 1899. 7

Figure 1.5. A toddler exhibits a highly variable gait while learning to walk. 8

Figure 1.6. Key structures of the primate motor system. 10

Figure 1.7. The motor homunculus as drawn by Wilder Penfield, 1950. 12

Figure 1.8. Regions with projections to dorsal (yellow) and ventral (blue) PMC. 15

Figure 1.9. Representation of sequence features in PMC revealed using multivariate pattern analysis of fMRI data. 17

Figure 1.10. Retrograde tracing from SMA and pre-SMA in rhesus macaque show dissociation of input structures. 19

Figure 1.11. Example neuron responses in pre-SMA and SMA during sequence performance. 20

Figure 1.12. The basal ganglia. 21

Figure 1.13.. The corticobasal-ganglia thalamocortical loops. 22

Figure 1.14. Dendritic spine formation in mouse motor cortex after learning a novel reaching movement. 27

Figure 1.15. Performing an action requires activating a set of neurons in a population. 28

Figure 1.16. Stages of population tuning. 30

Figure 1.17. Learning by neural reassociation. 35

Figure 1.18. Diagrammatic representation showing the difference between the null space and potent space. 37

Figure 2.1. A prototypical implementation of the Serial Reaction Time Task. 45

Figure 2.2. Trial structure in the Force Tracking Task. 47

Figure 2.3. The magnetic moment. ____ 49

Figure 2.4. Magnetic dipole moments exposed to a strong magnetic field. ____ 50

Figure 2.5. Precession and the Larmor Frequency. ____ 51

Figure 2.6. Nuclear magnetic resonance. ____ 52

Figure 2.7. Time course of the hemodynamic response. ____ 54

Figure 2.8. Demonstrating the feasibility of BOLD in humans. ____ 55

Figure 2.9. Blood oxygenation level closely matches spontaneous neural activity in rodents. ____ 57

Figure 2.10. Motion induced global artefacts in fMRI data. ____ 59

Figure 2.11. Multi-echo EPI improves SNR. ____ 61

Figure 2.12. Characteristic BOLD-like and noise-like components. ____ 66

Figure 2.13. High-resolution MRSI acquired at 3-Tesla. ____ 68

Figure 2.14. High-resolution metabolite mapping of a high-grade oligodendroglioma in a 52-year-old female. ____ 69

Figure 2.15. K-space sampling with concentric ring trajectory. ____ 72

Figure 2.16. Spatial response function (SRF) across k-space sampling methods . ____ 73

Figure 2.17. Spectral SNR of MRSI measurements from a phantom. ____ 74

Figure 2.18. Diagram of non-water suppressed MRS acquisition. ____ 75

Figure 2.19. Example basis set and LCModel fit from a single MRSI voxel. ____ 77

Figure 3.1. Experimental design. ____ 83

Figure 3.2. Improvement in behavioral performance after training. ____ 96

Figure 3.3. Training-specific changes in resting-state functional connectivity with the hippocampus. ____ 98

Figure 3.4. Training-specific changes in resting-state functional connectivity with the sensorimotor network. ____ 100

Figure 3.5. The magnitude of the interaction between Time-of-Day and Visit on resting-state functional connectivity of the hippocampus. ____ 102

Figure 3.6. The magnitude of the interaction between Time-of-Day and Visit on resting-state functional connectivity of the sensorimotor cortex. ____ 103

Figure 3.7. The relationship between performance improvement in the visuo-spatial task and hippocampal resting-state functional connectivity change during Visit 2. ____ 105

Figure 3.8. The relationship between short-term motor-sequence learning and sensorimotor cortex resting-state functional connectivity change during Visit 3. ____ *106*

Figure 3.9. Non-specific changes in resting-state functional connectivity observed due to time-of-day for hippocampus and sensorimotor cortex. ____ *108*

*Figure 4.1. Experimental design.*____ *123*

*Figure 4.2. Punishment improves performance of the SRTT.*____ *132*

Figure 4.3. Error-rate during the SRTT. ____ *136*

Figure 4.4. Punishment impairs to performance on the FTT. ____ *139*

*Figure 5.1. Experimental design.*____ *149*

*Figure 5.2. Punishment improves unintentional sequence learning.*____ *159*

Figure 5.3. Feedback did not impact awareness when learning unintentionally. ____ *162*

*Figure 5.4. Feedback does not impact performance during intentional sequence learning.*____ *164*

Figure 5.5. There was no difference between the feedback groups with respect to explicit knowledge formed during intentional sequence learning. ____ *166*

Figure 5.6. Comparison of early learning across the Feedback groups in both experiments. ____ *168*

Figure 6.1. Experimental design and skill memory retention. ____ *178*

*Figure 6.2. Reward and punishment differentially effect PMC functional connectivity change after training on the SRTT.*____ *183*

Figure 6.3. Punishment promotes PMC-striatal connectivity after FTT training. ____ *185*

*Figure 6.4. Regions showing PMC functional connectivity change after training had limited overlap across tasks.*____ *188*

Figure 6.5. Difference in functional connectivity change after training with feedback depends on task. ____ *189*

*Figure 7.1. Experimental design.*____ *201*

Figure 7.2. Example participant's aligned MRSI slabs. ____ *207*

Figure 7.3. Three sensorimotor tasks were performed to define the sensorimotor score. ____ *209*

*Figure 7.4. Correlation between the SNR+GM correction (regressing out the effects of SNR, grey matter concentration, and participant age) and the novel Cramer-Rao Lower Bound (CLRB) correction for both GABA and Glx.*____ *214*

*Figure 7.5. Metabolite distribution for GABA, Glx, tNAA, and tCho for 5 example participants.*____ *218*

Figure 7.6. CRLB values for GABA, Glx, tNAA, and tCho corresponding to participants in Figure 7.5. ____ *219*

Figure 7.7. SNR and LW maps for participants shown in Figure 7.5 and Figure 7.6. _____ *220*

*Figure 7.8. Map from each of the five representative scans for a single subject, along with representative spectra from nine contiguous, representative voxels.*_____ *220*

Figure 7.9. Distribution of reproducibility assessment measurements calculated in the present study. _____ *222*

Figure 7.10 Parcel-wise GABA and Glx concentrations in each Desikan atlas parcel with adequate coverage. _____ *224*

Figure 7.11. Correlation between concentrations of GABA and Glx across parcels. _____ *224*

Figure 7.12. Correlation between right superior frontal gyrus GABA concentration and Sensorimotor score. _____ *226*

*Figure 7.13 Correlation between GABA concentration in each Desikan atlas parcel with the behavioral metrics from each task.*_____ *227*

Figure 7.14. Correlation between GABA and functional connectivity strength for the left attention network. _____ *229*

Figure 7.15. The interaction between left attention network GABA concentration and mean network-connectivity strength relates to sensorimotor score. _____ *230*

List of tables

Table 2.1. Preprocessing pipelines used for fMRI studies. _____ *62*

Table 3.1. Inclusion of individual scans based on motion estimates. _____ *92*

Table 6.1. Results of the limited mixed effects model analysis for the SRTT and FTT. _____ *186*

Table 7.1. Quantification of data quality and reproducibility across participants. _____ *222*

Publications arising and supplementary work

The following publications and pre-print publications have arisen from the work in this thesis:

Adam Steel, Mark Chiew, Peter Jezzard, Natalie Voets, M. Albert Thomas, Charlotte Stagg, and Uzay E. Emir. New metabolite cycled density-weighted concentric rings k-space trajectory (DW-CRT) enables 1H magnetic resonance spectroscopic imaging at 3 Tesla. Scientific Reports, 05/2018; 8:7792. DOI: 10.1038/s41598-018-26096-y ***[Chapter 2]***

*Adam Steel *, Cibu Thomas*, Aaron Trefler, Gang Chen, Chris I. Baker (2019). Task-specific changes in intrinsic functional connectivity accounts for the evolution of motor and visuospatial skill learning). Neuroimage, 118:524-538. DOI: https://doi.org/10.1016/j.neuroimage.2018.12.038.* ***[Chapter 3]***

Adam Steel, Edward H. Silson, Charlotte J. Stagg, Chris I. Baker (2016). The impact of reward and punishment on skill learning depends on task demands. Scientific Reports, 10/2016; 6:36056. DOI: 10.1038/srep36056. ***[Chapter 4]***

*Adam Steel, Edward H. Silson, Charlotte J. Stagg, Chris I. Baker (2019). Differential impact of reward and punishment on functional connectivity after skill learning. Neuroimage , 189:95-105. DOI: https://doi.org/10.1016/j.neuroimage.2019.01.009 [****Chapter 6]***

In addition, I contributed to the following works, which were outside the scope of this thesis and therefore not included:

Tim Lawrence, Adam Steel, Martynn Ezra, Mhairi Speirs, Pieter Pretorius, Gwenaelle Douaud, Stamatios Sotiropoulos, Thomas Cadoux-Hudson, Uzay Emir, Natalie Voets (2019). MRS and DTI evidence of progressive posterior cingulate cortex and corpus callosum injury in the hyper acute phase after TBI. Brain Injury. DOI: https://doi.org/10.1080/02699052.2019.1584332.

Edward H. Silson, Adrian W. Gilmore*, Sarah E. Kalinowski, Adam Steel , Alexis Kidder, Alex Martin, and Chris I. Baker (2018). A posterior-anterior distinction between scene perception and scene construction in human medial parietal cortex. The Journal of Neuroscience , 1219-18; DOI: https://doi.org/10.1523/JNEUROSCI.1219-18.2018 .*

Edward H. Silson, Adam D. Steel, & Chris I. Baker (2016). Scene-Selectivity and Retinotopy in Medial Parietal Cortex. Frontiers in Human Neuroscience, 08/2016;10(412). DOI: 10.3389/fnhum.2016.00412

Marcie King, Iris Groen, Adam Steel, Dwight Kravitz, Chris I. Baker. Similarity judgments and cortical visual responses reflect different properties of object and scene categories in naturalistic images. Submitted. Available on bioRxiv: https://www.biorxiv.org/content/early/2018/05/10/316554.article-info

Adam Steel, Cibu Thomas, Chris I. Baker (2018). Effect of time of day on reward circuitry. A discussion of Byrne et al. 2017. arXiv, arXiv:1802.00753 [q-bio.NC]. http://arxiv.org/abs/1802.00753

Mark Chiew, Wenwen Jiang, Brian Burns, Peder Larson, Adam Steel, Peter Jezzard, M. Albert Thomas, and Uzay E. Emir (2017). Density-Weighted Concentric Rings K-Space Trajectory for 1H Magnetic Resonance Spectroscopic Imaging at 7 Tesla. NMR in Biomedicine, 2018:31:e3838. DOI: 10.1002/nbm.38388

Acknowledgements

My thesis would not have been possible without the support of my mentors, labmates, friends, and family. I extend my sincerest thanks to everyone who has seen me through this fantastic journey. Below is a short list of those to whom I owe special thanks.

To my two mentors, Chris and Charlie. I could not have wished for more supportive people to shepherd me toward this goal.

Thanks to Chris for always pushing me to be a more thoughtful scientist and for flying me around the world and back. My time in your lab was nothing short of a party in the USA.

To Charlie for letting me jump into rabbit holes and also for pulling me out of them. I benefited tremendously from your taste in all things neuro- and vino-logical.

To the NIH-OxCam Scholar program and the Biomedical Research Alliance for their financial and administrative support.

To Ed Silson for giving up his nights, mornings, and weekends to help me acquire the data for Chapters 4 and 6. “They share a lot” doesn’t quite capture it. Thanks for being a true friend.

To Cibu Thomas for designing the experiment and collecting the data for Chapter 3. Dropping by your office never failed to lift my spirits. Thanks for helping me find the baby.

To Uzay Emir, who developed the MRSI sequence used in Chapter 7, for thoughtful guidance.

To Piergiorgio Salvan and Jacob Levenstein, for help preprocessing the resting-state fMRI data for Chapter 7. When you’re with your friends, it never feels like work.

To the PNG and Plasticity groups for karaoke and pub nights. Special thanks to Emily and Catharina for keeping me connected when I was in the US.

To Georgina James for helping collect data in Chapter 7.

To the SLaP Lab for memes, acceptance schnapps, and AlphaBears.

To Anna Marcus and Tyler Blake for sharing coffee and good cheer.

To Glenn Cahilly-Bretzin for 20-minute adventures and Fish Fridays.

To Michael Niklaus for always saying yes to “630 CCG?”

To Craig Schwait for sage advice.

To Morgan Steel for encouragement.

And to my parents, Gary and Linda Steel, for your endless support of me and my dreams.

To my wife, Lindsay.

Examining functional plasticity in the human brain

Adam Steel

Trinity term 2018 Lincoln College

A thesis submitted in partial fulfilment of the requirements for the degree of

Doctor of Philosophy in the University of Oxford

Abstract

The ability to quickly learn and adapt behaviour is critical to animal survival. Recent evidence has shown that short-term training can induce widespread changes in the neural activity across the brain. In this thesis, I conducted a series of studies that examined the behavioural and neural consequences of short-term training in humans.

In Chapter 3, I showed that resting-state functional connectivity could detect task-specific changes. However, I also observed changes that were not task-specific, and even occurred when participants had no training. This experiment demonstrates that resting-state functional connectivity is an appropriate tool for studying plasticity and also highlights the importance of experimental controls in plasticity research.

In Chapter 4 and 5, I investigated the impact of reward and punishment on performance during skill learning. In Chapter 4, I focused on how task-demands interact with skill the learning, and in Chapter 5, I focused on whether the learner's intention to learn modifies the impact of reward and punishment.

In Chapter 6, I used resting-state fMRI to elucidate the impact of reward and punishment on the neural activity evoked by training on two different tasks. I determined that reward and punishment evoked different neural patterns after training, but that the precise regions affected by feedback processing depended on the task being trained.

In Chapter 7, I conducted a preliminary investigation into the neurochemical basis of resting-state functional connectivity using a novel MRSI technique. I found a preliminary link between the distribution of the inhibitory neurotransmitter GABA and network-level functional connectivity. I also found that the interaction between GABA concentration and network-level functional connectivity may be associated with sensorimotor performance.

Collectively, these studies demonstrate the short-term training induces wide-spread changes throughout the brain. These changes are influenced by the task being trained as well as contextual influences. In addition, these findings raise new questions into the physiological basis of network level plasticity that can be addressed in the future.

This thesis contains approximately 48,000 words.

"I believe that we are subject to the law of habit in consequence of the fact that we have bodies….

Our nervous systems have … grown to the way in which they have been exercised, just as a sheet of paper or a coat, once creased or folded, tends to fall forever afterward into the same identical folds."

- William James, The Principles of Psychology, 1899.

1 Introduction

In this section, I give background information relevant to this thesis. I begin by defining skill learning, and subsequently give a brief overview of several brain areas involved in motor skill learning. Next, I discuss biological mechanisms for learning. I conclude the chapter by outlining the major aims of this thesis.

1.1 Skill: a curious thing

"With tennis, you can go pick up a racket, take a lesson, and understand how much talent and skill it takes to be as good as the top pros. Same with golf: pick up a club. But not many can go out and get in a race car and experience a drive over 200 miles an hour."

– Mario Andretti, Super Speedway (film), 1997



Figure 1.1. The Marina Bay Street Circuit, looking south toward the Marina Bay Sands hotel. Source: Wikimedia commons: 1_singapore_f1_night_race_2012_city_skyline.jpg

On a steamy September night, a Formula One car hurdles through the streets of Singapore (Figure 1.1). Traveling 250 km/h, the driver traverses Anderson Bridge, screams past the Fullerton Hotel, and brakes hard to make a sharp turn onto Esplande Drive beside Merlion Park. With deft coordination of the gas, brake, and gearing, the driver negotiates four consecutive turns, left, right, right, left, in quick succession. Then, to the delight of the cheering spectators, the engine roars as the driver takes the final left turn and accelerates to over 300 km/h, passes the Singapore Flyer, and takes the flag to claim victory in the Marina Bay Street Circuit, one of the most spectacular racing events in the world.

Defining 'skill' is almost impossible. Undoubtedly, driving a Formula One vehicle requires tremendous talent and years of practice. The high speeds impart immense forces on a driver's body and pushes the ability to process visual information to its limit. The sensory information must be transformed into motor commands for the arms and hands to steer and to the feet to work the pedals. Yet this is only one example of a 'skill.' A basketball player's three-point shot, an archer's draw, or a dancer's step are all examples of skill, and, objectively, these actions have very little in common. Yet, despite their dissimilarity, classifying each of them as skills is not controversial. This makes studying the acquisition of skills highly complex – how can we study something that we cannot easily categorise?

Given the diversity of behaviours that might be considered skills, one might throw up their hands and say, "You know it when you see it" (YKIWYSI). But clearly this definition is insufficient.

YKIWYSI leaves open the possibility that skill is conflated with an innate capacity for performance (talent). 'Raw talent' is a common expression to describe an athlete who, despite not practicing, may perform very well. But, raw talent

cannot explain some highly specific skills developed. For example, classical archers may become disproportionately adept at hitting their target from the competition distance (18m), which cannot be explained by raw talent alone (Figure 1.2) [1]. Therefore, it is important to disambiguate performance ability from innate talent from practice when defining skill.

How then to define the inscrutable 'skill'?

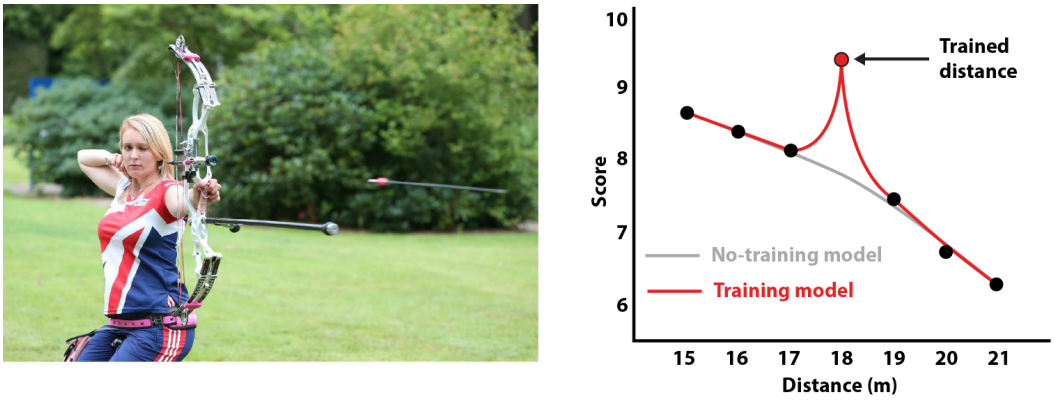


Figure 1.2. Classical archers exhibit a highly specific skill. Left. An archer from the Great Britain national team. Right. In a study by Nabavanik and colleagues, classical archers, who shoot from 18 meters during competition, were asked to take 10 shots from 15-21 meters. Performance was significantly better at the competition distance, which is practiced most often. Grey curve shows a quadratic function fitted to the mean scores at 15, 16, 17, 19, 20, and 21 m (black points), excluding 18 m (red point). The red curve shows the fit of a model that incorporates both distance and practice as factors. Left photo taken from ***https://www.bristolmedia.co.uk/news/1635/archery-gb-aims-for-pr-gold-with-mccann-bristol****. Right figure redrawn from [1].*

1.1.1 What defines skilled behaviour?

Generally, three features distinguish unskilled and skilled performance: 1) goal completion, 2) reproducibility, and 3) efficiency. Goal completion is often considered so obvious that it is sometimes only considered implicitly when defining skilled behaviour [2-4]). However, I include completion in my definition

because the goal of an individual is critical when considering behavioural performance.

The second point I consider necessary is reproducibility. Reproducibility is the ability to consistently produce a behaviour that obtains the goal in the same way. Reproducibility is generally achieved through repetition. Through practice, the behaviour goes from being highly-variable initially to converge on a stable, 'final' form as the behaviour's defining qualities become more similar for each repetition.

The third point I consider necessary for skill is efficiency. Efficiency refers to the balance between the positive outcome after success compared to the costs of the behaviour's performance. For many behaviours, the energy required for performance is a critical cost, while some external reward (e.g., food) is the benefit of performance. However, external costs can also play a role in shaping behavioural performance. Likewise, there may be an inherently rewarding quality to completing a goal. All of these factors must be considered when defining skill.

Taking together the importance of completion, reproducibility, and efficiency in the definition of skill, in this thesis I define 'skill learning' as: The convergence of a behaviour, through repetition, to a stereotypical form optimized for the task goal given the constraints under which it was repeated.

1.1.2 How to measure skill learning?

Complex skills, like driving, are composed of behaviors: extracting information from road signs while traveling at high speed, coordinating the gas and brake pedals, and adjusting the radio. Likewise, driving tests query specific relevant behaviors: drivers have a vision test, execute a three-point turn, and parallel park. In the lab, skill is often assessed within a *domain*, or an approximate

categorization of its requirements. For example, *visual-perceptual skill* might be assessed using a texture discrimination task, where the participant detects the location of three diagonal lines presented for a very short duration. Improvement in texture discrimination would be considered a type of *perceptual skill learning* [5] (Figure 1.3). Similar tasks have been developed to assess perceptual skill for audition [6, 7]. Other domains include *motor skill [8]* and more abstract abilities, such as *object skill* for selecting high-value objects [9].

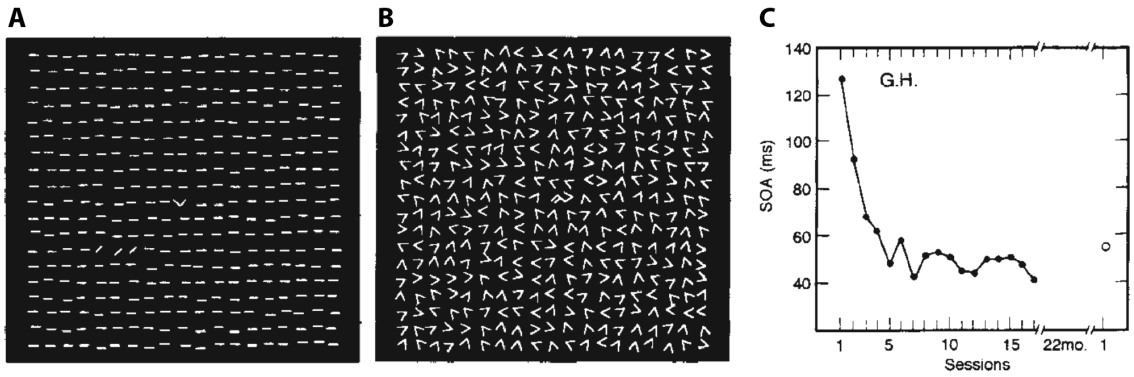


Figure 1.3. Perceptual skill learning assessed with a texture discrimination task. A. The target stimulus (three diagonally oriented lines) is rapidly presented (e.g., 10 ms duration) to the participant, followed by a masking stimulus after some duration (B). Performance on the task is determined by the duration between the presentation of the target and mask stimuli, where the participant achieves 95% correct responses (stimulus onset asynchrony). C. Representative participant learning curve. Across sessions, the stimulus onset asynchrony where the participant achieves 95% criterion is reduced. In addition, the skill is maintained for 22 months despite not practicing the skill. Figure adapted from [5].

1.1.3 Motor skill learning

In this thesis, I primarily consider *motor skill learning* (Chapter 3 also considers visuo-spatial learning). The repositioning of an aspect of the body in space, *movement,* is a grounding principle of motor skill. Thus, I define *motor skill* as *a movement that has reached a stereotyped form that is optimized to achieve a*

goal, while maximizing the possible external reward and minimizing energy expenditure as well as external costs.

1.1.3.1 Measuring motor skill

The success of a movement is usually measured in spatial (how far did the movement deviate from the goal?) and temporal (did the movement occur at the desired time?) terms, and these components are often considered to be related. Their relatedness is particularly apparent when the speed of execution is prioritized: faster movements generally produce less accurate outcomes. This phenomenon is generally referred to as the speed-accuracy trade-off [10, 11]. In a series of seminal experiments for his doctoral work at Columbia University, Robert Woodworth demonstrated the speed-accuracy trade-off in movement by having participants draw lines to a precise distance using a kymograph, a tool for drawing straight lines, while keeping time with a metronome (Figure 1.4). Woodworth observed that as he increased the metronome's rate, the participants drew lines that deviated more from the target length.

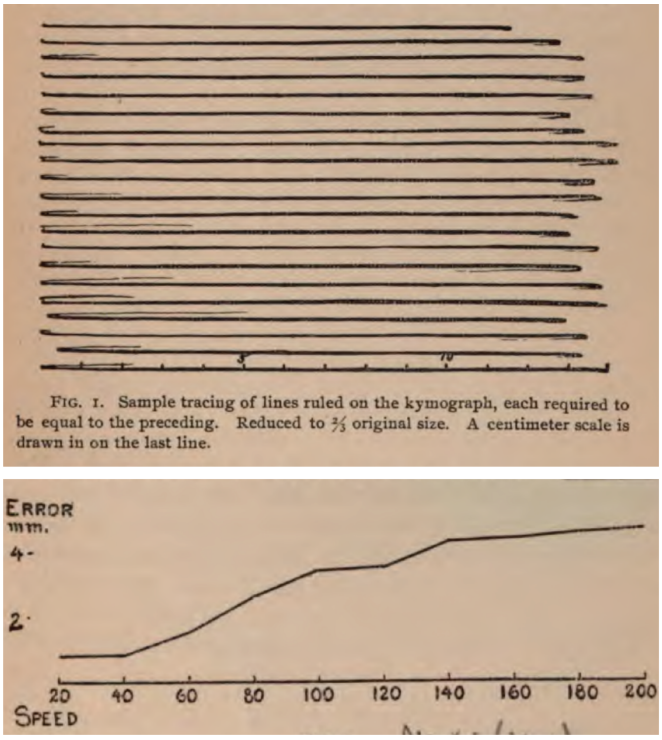


Figure 1.4. The speed-accuracy tradeoff demonstrated by Woodworth, 1899. Upper. Sample kymograph tracings by a single participant. Participants were instructed to draw lines of equal length keeping time with a metronome. Lower. As speed of performance increases, the average error in line length increases. This effect has been replicated in many modalities. Figures reproduced from [11].

The speed-accuracy trade-off is apparent for other types of hand movements [3], but also in less obvious movements like mouth and vocal cord movements during speech [2]. The combination of speed and accuracy can be a proxy for the efficiency of a movement, and prior work has operationalized skill as the slope of a psychometric function fit to the speed-accuracy trade-off of participant's movements [3, 4]. A generally acceptable operationalization of skill learning is a 'shift' of the speed-accuracy trade-off, which may manifest as an increase in accuracy at a given speed, an increase in speed without sacrificing accuracy, or both.

For movements where the primary goal is speed, showing that accuracy remains high is particularly important. However, using the speed-accuracy trade-off as *the* definition for motor skill can neglect a critical feature of skill defined

above: reproducibility of the behaviour, or the reduction in variability. The insufficiency of speed-accuracy trade-off as a definition for motor skill is most notable when considering movements that are not optimized to achieve greater speed, such as walking [12].

The extent to which walking is a skill *per se* is somewhat debated, but the development of this ability follows skill-like trajectory. Early in development, children exhibit highly variable gait patterns during walking (Figure 1.5). As children continue walking, the gait pattern stabilizes to one that may persist into adulthood and destabilizes in old age [13, 14]. Notably, although the neural mechanisms for learning during childhood may differ from adulthood in some respects, the behavioural outcomes being optimized (economy of movement, smoothness, energy efficiency) rather than speed are present throughout the lifespan.



Figure 1.5. A toddler exhibits a highly variable gait while learning to walk. The toddler shown here is receiving assistance from a highly-skilled adult.

In the laboratory setting, force tracking experiments, where participants must make movements to follow a smoothly moving target, can be used to measure motor performance agnostic of speed [15-18]. In these experiments, the outcome measure is the reduction in tracking error, which generally results from a reduction in movement variability, during movement production.

1.1.3.2 Summary

In summary, motor skill learning is the optimization of movement patterns to achieve a goal that occurs through practice. Depending on the goal of the behaviour, skill learning can be operationalized as a shift in the speed-accuracy trade-off or as a reduction in variability of a movement. In the next section, I discuss a few important brain regions (Figure 1.6) involved in movement execution and their potential involvement in motor skill learning.

1.2 The motor system

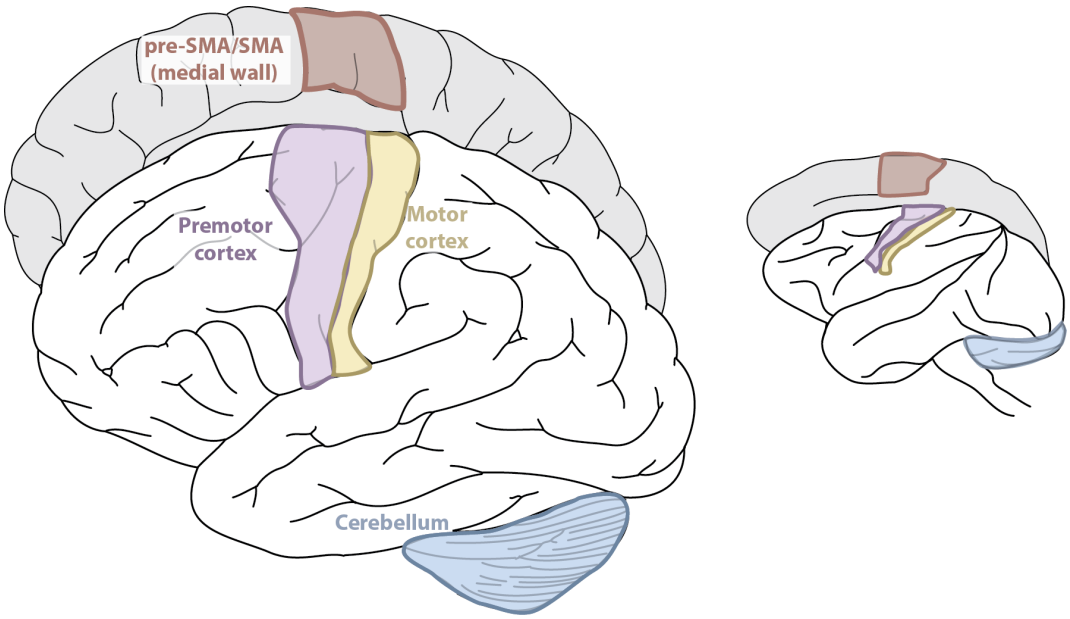


Figure 1.6. Key structures of the primate motor system. Both the human (left) and macaque (right) are shown. The medial wall of the right hemisphere is shown in gray. include the primary motor cortex (yellow), premotor cortex (purple), and the supplementary motor area (brown, drawn on the medial wall of the right hemisphere). Although not part of the motor system proper, the cerebellum (blue) and basal ganglia (not shown, see Figure 1.12) play critical roles in motor learning. Figure is drawn approximate to scale.

Movement is crucial to the survival of most terrestrial animals, and, perhaps driven by this necessity, most mammals have developed a cortical area that projects directly to the spinal cord, usually referred to as the primary motor cortex or M1. However, primates, particularly humans, are unmatched in our ability to dexterously control our forelimbs. Concurrent with this ability, other specialized cortical areas devoted to the spatial and temporal aspects of movement planning (premotor cortex [PMC] and supplementary motor area [SMA], respectively), have increased in areal extent. In this section, I discuss role of each area in motor control. I also discuss the role of the basal ganglia and

cerebellum, which provide signals for task success and error correction, respectively.

1.2.1 Primary motor cortex

The primary motor cortex, situated primarily on the precentral gyrus is the main output area from the brain to the spinal cord for non-reflexive muscle [19].

At a coarse level, the spatial topology of motor cortex broadly follows that of the human body, a feature known as somatotopy (although the organization is more complex at a finer level, see below). This coarse mapping between the motor cortex and the body is often represented as the motor homunculus, drawn by Wilder Penfield based on the muscle activity evoked by direct electrical stimulation of the brain during surgery (Figure 1.7). Rather than represent the body in proportions that reflect real-world size, the areal extent of each body part is broadly dictated by the extent to which it can be dexterously controlled; thus, the face and hands are represented over a large cortical area, while the trunk is represented over a relatively small area.

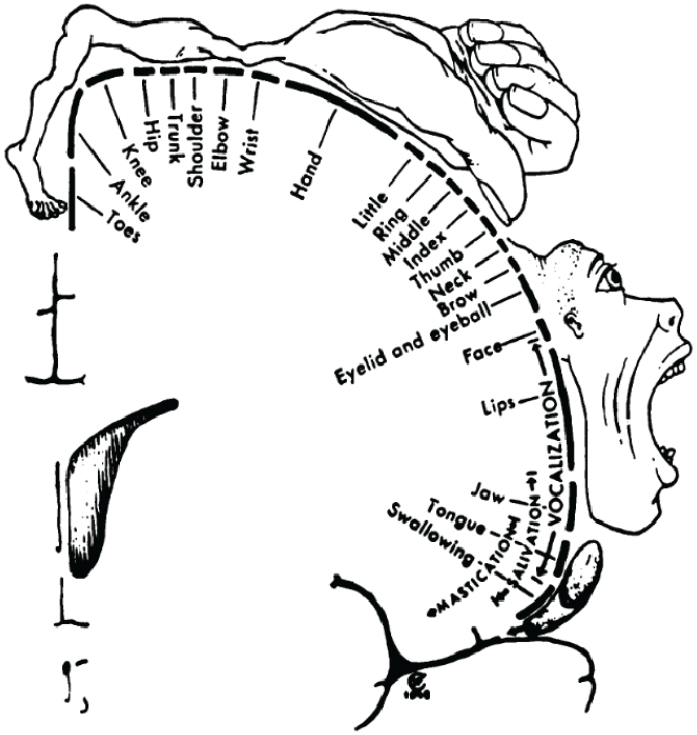


Figure 1.7. The motor homunculus as drawn by Wilder Penfield, 1950. This drawing depicts the locations along the precentral gyrus where electrical stimulation elicited overt movements. The size of the body part shows the approximate proportion of cortex devoted to that body part, emphasizing the over-representation of the hands, face, and mouth compared to other body parts. Figure reproduced from [20].

One possible organization for muscular control is a pure somatotopic arrangement with a specific area-to-muscle mapping, where a discrete set of neurons might represent the flexion of a given muscle. However, more recent electrophysiological recordings in non-human primates suggest that this is not the case [21]. For example, if primary motor cortex were laid out in a purely somatotopic fashion, naturalistic bouts of electrical stimulation (long trains, e.g., 500 ms) in a single location should evoke persistent contraction of a single muscle. However, naturalistic stimulation actually evokes complex and coherent movements [22, 23]. Indeed, when investigated at a more fine-grained level, the somatopic organization of primary motor cortex appears far more convoluted [21,

24]. Stimulation within the overall representation of the arm reveals that specific joint movements, such as extension of the wrist or elbow, may be represented in multiple locations. In addition, the cortical representation of neighbouring joints may overlap with one another [24-30].

The evidence presented above suggest that primary motor cortex is not laid out in a purely somatotopic fashion. Instead, some have speculated that motor cortex organization reflects an 'ethological' map [23], where behaviours are grouped by type (i.e., reaches are grouped together, defensive behaviours are grouped together, etc.). Others have suggested that motor cortex is organized by muscle synergies or 'natural use,' where the layout of motor cortex reflects movements that commonly occur together, but not necessarily by behavioural phenotype [27, 31]. While there is not yet consensus on the nature of the representation, importantly for this thesis, these representations appear to be malleable [32, 33]. The reorganization of cortical representations from skill learning is explored in more detail in Section 1.3.3.

1.2.2 Premotor cortex

The PMC, immediately anterior of primary motor cortex, is thought to play an integral role in movement planning, particularly in the transformation of sensory signals into motor actions [34-36].

Studies using primates with PMC lesions provided evidence that the PMC is causally involved in sensory guided movement [37]. To facilitate visually guided movement, the PMC receives input from the dorsal visual stream via connections with the intraparietal sulcus (IPS; [38]). In IPS, neurons encode the spatial location of objects in the visual field in a spatial-invariant manner, i.e., real-world position relative to the body [39]. In turn, neurons in the PMC, remap these locations in an

effector-centric manner, i.e., the position of the object to be grasped relative to the hand [40]. In addition, to facilitate grasping, as well as the target position, PMC neurons also represent the relative position of the hand and eyes [41].

With respect to organization of this area, anatomical and electrophysiological evidence suggest that the dorsal and ventral aspects of PMC have different functions for movement planning. Each has distinct anatomical connections with IPS: ventral and dorsal PMC have anatomical connections with the ventral IPS and anterior IPS, respectively. These subdivisions of IPS are thought to represent specific information about the target stimulus. During grasping, neurons in ventral IPS respond to spatial properties of the target, such as motion [42], while neurons in anterior IPS represent structural properties of the object necessary to act upon it, such as orientation and size [43, 44]. These inputs may help facilitate their overall function (see below).

In addition to input from ventral IPS, the macaque ventral PMC also receives input from rostral prefrontal cortex. In contrast, dorsal PMC does not have connections to prefrontal cortex [45] (Figure 1.8). Tractography-based parcellations of PMC using diffusion weighted imaging suggest that these differences persist in humans [46].

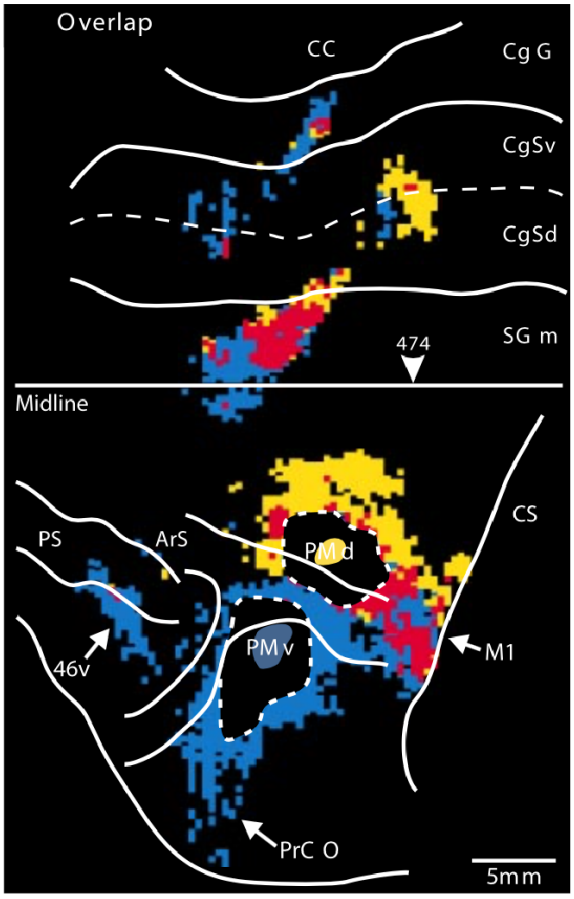


Figure 1.8. Regions with projections to dorsal (yellow) and ventral (blue) PMC. Retrograde tracer injections reveal that area 46v, in rostral prefrontal cortex, projects to ventral, but not dorsal, PMC. In addition, primary motor cortex projects to both regions. Only regions anterior of the central sulcus are shown. Figure adapted from [45].

The anatomical connectivity of the ventral and dorsal PMC suggests that they may be specialized for 'long-term' motor planning (e.g., planning the trajectory of a reach or selecting an object to reach for) and 'short-term' motor planning (e.g., how to interact with an object of a specific shape), respectively. The functional properties of dorsal and ventral PMC corroborate this account. Ventral PMC neurons primarily respond to more distant targets. In contrast, neurons in dorsal PMC respond more strongly to proximal targets, and also represent the action required by the target [36, 40, 47-50].

Dorsal PMC may be particularly important for selecting among actions, and appears to represent the set of possible option prior to movement [51]. Indeed, the causal role of this region in an action selection task has also been shown using transcranial magnetic stimulation (TMS) [52]. In this study, participants were asked to make an appropriate response to a stimulus by integrating across two stimulus features (size and shape). After inhibitory TMS was applied to both premotor cortices, response time (RT) in trials that required integration across stimulus features increased. However, RTs on simple reaction time trials that did not require feature integration remained at pre-stimulation levels. This provides strong evidence for the role of bilateral dorsal PMC in integration of abstract stimulus features to select appropriate actions.

Evidence also suggests dorsal PMC plays a critical role in visuomotor skill learning and sequence learning. Primates with PMC lesions fail to learn and recall arbitrary visual-motor associations, such as a hand movement, in response to a stimulus on a computer screen [53]. In addition, when learning arbitrary visuomotor associations, dorsal PMC neurons develop stimulus-specific responses [34, 54]. Learning visuomotor tasks, including sequences of finger movements, reliably evokes PMC blood oxygen level dependent (BOLD)-signal changes measured with functional magnetic resonance imaging (fMRI) [55]. Work by Weistler and Diedrichsen as well as Kornysheva and colleagues suggest that activity changes in PMC may reflect the development of movement sequence representations in PMC [56, 57]. After learning has occurred, aspects of the learned sequence could be 'decoded' from the BOLD-response in PMC. The decoded information included the temporal and spatial components of learned sequences (Figure 1.9).

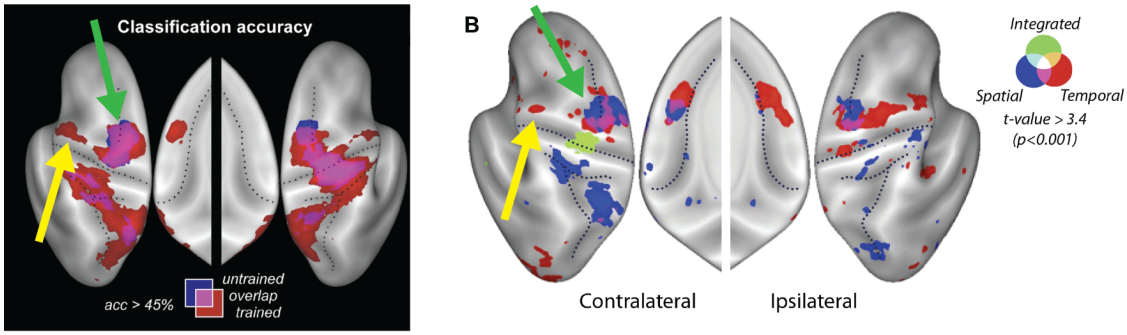


Figure 1.9. Representation of sequence features in PMC revealed using multivariate pattern analysis of fMRI data. A. Weistler and Diedrichsen trained participants to perform either four previously learned (trained) or four novel (untrained) sequences of button presses using the left hand. The authors then looked for regions where the pattern of the BOLD-response among neighboring voxels could be used to distinguish between the trained or untrained sequences (decoded). This process revealed that both trained and untrained sequences could be decoded from dorsal PMC (green arrow), while only trained sequences could be decoded from ventral PMC (yellow arrow). B. Using a similar procedure, Kornysheva and colleagues looked for regions where sequences of trained movements could be classified based on their spatial or temporal structure. During performance, both spatial and temporal information could be decoded from dorsal PMC, while only temporal information could be decoded from ventral PMC. Panel A is adapted from [57]. Panel B is adapted from [56].

1.2.3 Motor areas on the medial wall

Together, SMA and pre-SMA, located dorsal to the cingulate gyrus on the medial wall, organize movements into sequential order, and facilitate the initiation, or gating, of volitional movement [58].

The role of SMA and pre-SMA in volitional control of movements is potentially mediated through connections with the basal ganglia in a circuit specialized for response inhibition [59]. Electrical stimulation of the SMA produces complementary effects with respect to voluntary movements; stimulation of these regions while at rest induces complex, naturalistic movements of the face, arm, or leg. In contrast, stimulation of these regions during ongoing movements of the arms or during speech cause freezing or repetition of the ongoing movement [60]. In addition, patients with lesions in SMA demonstrate

the importance of this circuit for movement inhibition; these patients can develop *alien hand syndrome*, whereby the limb contralateral to the lesion executes complex movements without the conscious intent of the patient [58, 61].

In addition, SMA and pre-SMA play a critical role in the temporal organization of movement sequences. Monkeys with lesions to SMA fail to execute sequences of two or three saccades held in short-term memory despite having no problem making single, directed saccades, [62]. Temporary inhibition of SMA with transcranial magnetic stimulation produces a similar deficit in human participants when performing sequences of finger movements [63].

Both the structural connections and functional responses of SMA and pre-SMA suggest that these regions are hierarchically organized. Although they are densely interconnected, the SMA and pre-SMA receive differing input. SMA receives projections from more posterior motor structures on the lateral surface, including caudal PMC and primary motor cortex. In contrast, pre-SMA receives input from more anterior motor structures, including rostral portions of the premotor cortices and rostrolateral prefrontal cortex (Area 46v, which also projects to ventral PMC). Pre-SMA receives very little input from primary motor cortex (Figure 1.10) [64].

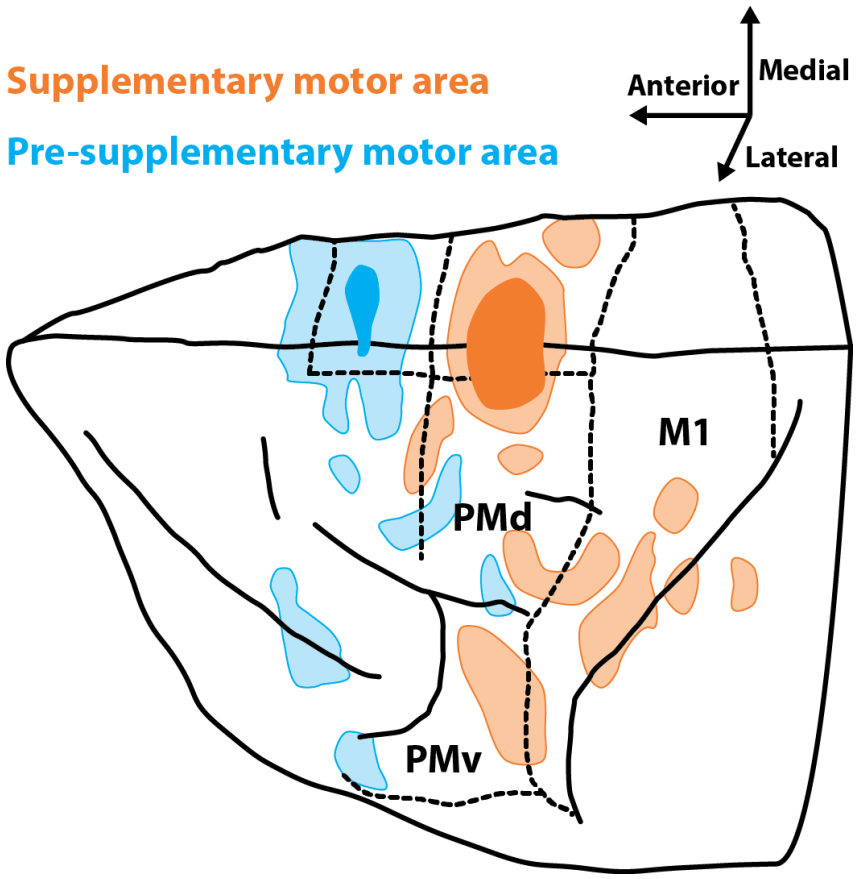


Figure 1.10. Retrograde tracing from SMA and pre-SMA in rhesus macaque show dissociation of input structures. SMA (blue) primarily receives projections from posterior motor cortical structures. Pre-SMA (orange) primarily receives input from anterior motor cortical structures, as well as lateral prefrontal cortical structures. Dark colours indicate injection locations. M1: primary motor cortex. PMd: dorsal premotor cortex. PMv: ventral premotor cortex. SMA: supplementary motor area. pre-SMA: Pre-supplementary motor area. Figure redrawn from [64].

Functionally, when performing a sequence of movements from memory, during the interval between movements, significant SMA neurons encode the *transitions between two movements* in a sequence. In contrast, pre-SMA neurons encode the *ordinal position within a sequence* (Figure 1.11) [65]. In addition to performance from memory, SMA and pre-SMA neurons exhibit different responses during learning [66]. A significant population of neurons in pre-SMA respond to new movement sequences, but the activity of these neurons *decreases* over the learning process. Once a sequence is well learned, very few pre-SMA neurons will respond prior to its execution. On the other hand, a larger proportion

of SMA neurons respond to well-learned sequences (although cells with other response properties constitute the majority of the population in this region, see [66]). FMRI studies in humans have also demonstrated the involvement of pre-SMA in the early sequence learning [67]. In addition, consistent with electrophysiological findings, sequence identity, including the spatial and temporal components of the movement sequence, can be decoded from SMA [56, 57].

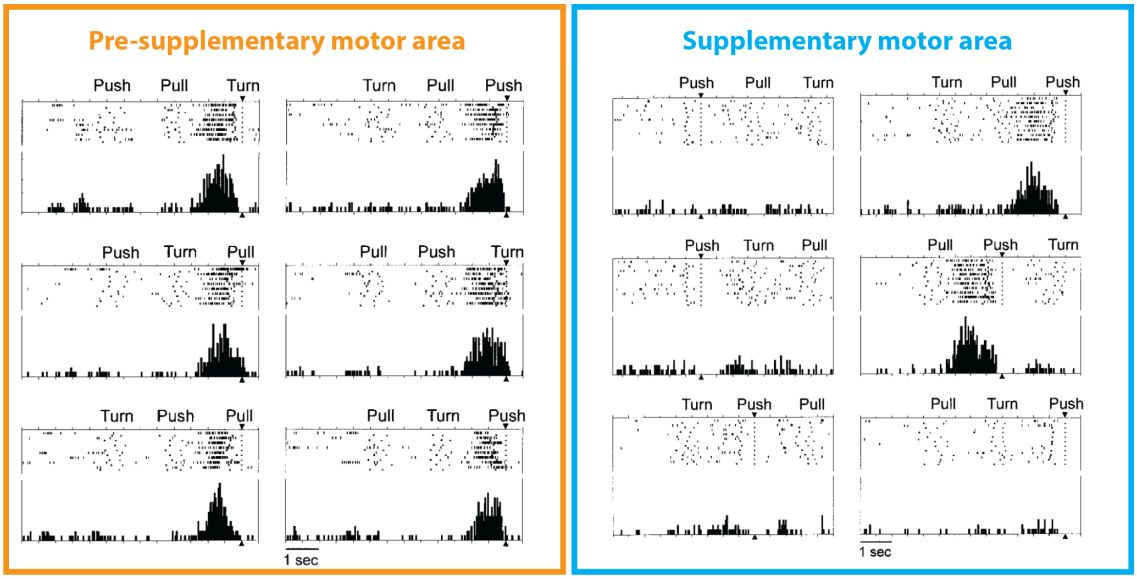


Figure 1.11. Example neuron responses in pre-SMA and SMA during sequence performance. Macaques were trained to perform sequences of Pull, Turn, and Push movements. Neuron responses were recorded from pre-SMA and SMA. Left. A neuron in pre-SMA showing selective response to the third movement, regardless of movement type. Right. A neuron in SMA showing selective response to the transition from Push to Pull movements. Figure adapted from [65].

1.2.4 The basal ganglia

To adapt behaviour, one strategy is to learn the value (likelihood of obtaining a positive outcome) of specific actions gradually through exposure [68]. This process is referred to as reinforcement learning, and the basal ganglia, a cluster of subcortical structures including the caudate, putamen, globus pallidus,

and subthalamic nucleus, are thought to implement valuation and reinforcement learning in the brain (Figure 1.12).

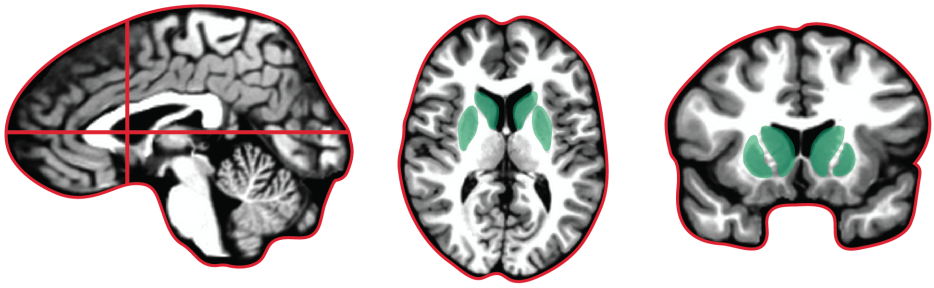


Figure 1.12. The basal ganglia. These structures (green) include the caudate nucleus, the putamen, and the globus pallidus.

The basal ganglia have broad connections throughout cortex, including with the orbitofrontal and limbic areas, prefrontal areas, sensory areas, and the cortical motor areas [69, 70] (Figure 1.13). These connections are generally described as a series of parallel, feedforward loops (corticobasal-ganglia thalamocortical loops) with four principal stages: 1) cortical regions project to the striatum; 2) striatal areas project to the pallidum; 3) pallidal projections then synapse in the thalamus, and 4) the thalamus projects back to the cortex [70-72]. More contemporary models of the corticobasal-ganglia thalamocortical loops include reciprocal connections between the stages of the loops as well as and horizontal connections amongst them, although I will only discuss a simplified model of this circuitry in this thesis [70, 73, 74].

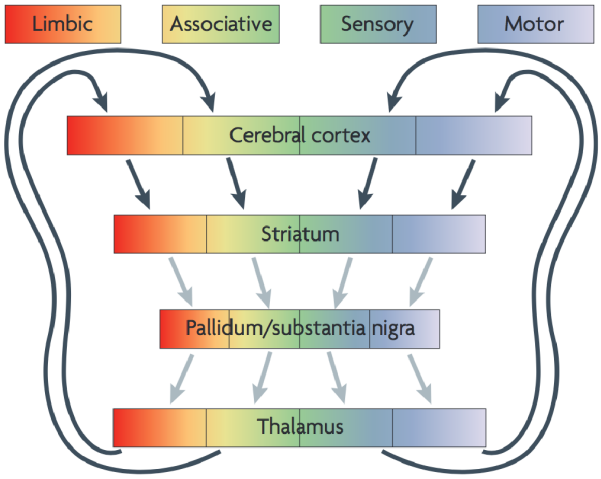


Figure 1.13.. The corticobasal-ganglia thalamocortical loops. Organization of the basal ganglia structures (caudate, putamen, and globus pallidus shown here) follow a gradient from limbic-associative to sensorimotor along their anterior-posterior axis. These connections play a key role in reinforcement learning. Figure reproduced from [70].

Most important for this thesis, the corticobasal-ganglia thalamocortical loops play a key role in the implementation of reinforcement learning to establish the value of an action in the human brain [75, 76]. In the context of motor behavior, action values established by 1) recursively selecting an action, 2) evaluating the outcome and 3) subsequently modulating the probability that the action will be selected in the future [75].

Initially, an action is selected among a set of possibilities. Whether the basal ganglia actively select the action based on cortical activity [77-79], or whether the action is selected at the cortical level informed by basal ganglia activity [76, 80] is debated [75]. However, it is generally agreed that the selection is based on the balance of activity between two cortico-basal ganglia-thalamocortical pathways, one excitatory and one inhibitory [71, 75]. Activity of the excitatory pathway selects (or excites) an action for execution, and the inhibitory pathway suppresses (or inhibits) competing action plans. A third pathway, which provides a global

stopping signal, also exists, but is not considered in the simplified model presented here [73].

Once an action is performed, dopaminergic signals arriving from the ventral tegmental area (VTA) provide a prediction-error signal [81-84]. The prediction error signals update the selection probability. Positive prediction errors, outcomes better than expected, result in a burst of neuronal activity in the VTA. The dopamine burst facilitates long-term potentiation (LTP) in the excitatory pathway, and long-term depression (LTD) in the inhibitory pathway; this joint modification leads to an increased likelihood of that action's selection in the future. An outcome worse than expected (a negative prediction error) results in a transient decrease in VTA neuron firing, which causes a decrease in dopamine concentration in the basal ganglia. The decrease in dopamine concentration induces LTD in the excitatory pathway and LTP in the inhibitory pathway and diminishes the likelihood that this action will be selected in the future.

1.2.5 The cerebellum

Although reinforcement learning strategies are effective, this learning strategy is often slow [68]. To enable rapid learning, before producing an action, one might build a model of a predicted outcome. After the action is complete, the actual outcome can be compared to the predicted outcome, and adjustments can be made based on the deviation. This process, forward-model comparison, is thought to be facilitated by the cerebellum [12, 85-87].

Like the basal ganglia, the cerebellum has connections to many parts of the cortex [88-90], and these connections are organized in a series of parallel loops. Input to the cerebellum comes from two primary pathways, 1) from cortex via mossy fibres originating in the pons, and 2) subcortical areas including the spinal

cord via climbing fibres originating in medulla [91]. These inputs provide the signals to enable forward-model comparison by conveying the system's estimated state to the actual state, respectively. An extremely reduced description of this circuitry is presented below.

The Purkinje cell, a large arboreous neuron found only in the cerebellum, is the workhorse for building the forward model and comparing the model and outcome. A simplified flow of the model-building and evaluation process might be the following: At movement onset, information about the current state of the cortex is sent via the mossy fibres, which synapse on the granule cells. Granule cells send long axonal projections across the cerebellar cortex and form synapses with many Purkinje cells. This distributes the cortical input widely across the cerebellum.

Because of their wide dendritic array, Purkinje cells can collect information from a large set of granule cells. As many as 20,000 granule cells may synapse on a single Purkinje cell. This can provide a comprehensive sampling sufficient to construct a model of the system's estimated state.

At the same time, input from the periphery arrives via the climbing fibres. In contrast with the granule cells, a single climbing fibre synapses directly onto the cell body of a single Purkinje cell. This provides a very specific mapping of the system's actual state among the Purkinje cells. The Purkinje cell integrates across these two input populations. The collective activity of the Purkinje cells is integrated in the deep cerebellar nuclei and relayed to cortex for behavioural adjustment.

By comparing motor actions to their sensory consequences, the cerebellum helps the motor system adjust to perturbations, such as the ability to adapt one's

gait to compensate for an uneven sidewalk [92]. Experimental research has demonstrated the cerebellum's importance in learning to compensate for rotated visual information (e.g., offsetting the visual information by some amount using prisms). For example, patients with cerebellar damage fail to compensate for a visual rotation when throwing a ball at a target [93]. Similarly, these patients fail to compensate for imposed force-fields while reaching [94].

1.2.6 Summary

The primate brain has evolved an elaborate system for motor control. These brain areas receive input from the sensory and prefrontal cortices to plan and execute movements. In addition, the motor control system receives input from subcortical structures to promote adaptive changes. The biological substrate for these adaptive changes is described below.

1.3 Neural mechanisms for learning

Mastery of a skill requires repeated performance, and, as Donald Hebb proposed in 1949, so too should its biological underpinning: *"when an axon of cell A is close enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth or process in both cells such that A's efficiency as one who excites B is increased"* [95]. The following section briefly discusses the Hebbian learning, or activity dependent plasticity, and the role of synaptic changes in the context of a complex system like the brain.

1.3.1 Activity dependent plasticity

Synapses are considered the most fundamental units for information transduction in the nervous system, and synaptic modification is thought to be the primary means for learning. Although non-Hebbian learning mechanisms, e.g.,

synaptic scaling [96-98], can modify neural response properties, LTP and LTD, the activity dependent processes for strengthening and weakening of synapses (respectively), are the most well studied [99].

In the first demonstration of LTP, Bliss and Lomo found that repeated electrical stimulation of neurons in the dentate gyrus of the hippocampus increased the response amplitude of a population of neurons downstream [100]. Later, Stanton and Sejnokski found that the timing of the downstream neurons' activation determined the nature of the activity dependent plasticity, i.e., whether LTP or LTD occurred [99, 101]. The metabolic cascade for Hebbian learning has been described in great detail, but a full description is beyond the scope this thesis [99, 102].

Activity-dependent plasticity leads to enlargement of the synaptic varicosities along the dendrites (spines), formation of new spines, or degeneration of spines [103-105]. In this thesis, I collectively refer to these processes as *synaptic modification*. These modifications have been observed throughout the brain, including the visual cortex [106], motor cortex [107], and the midbrain dopamine circuit [108].

1.3.2 Synaptic modification for learning

In 2009, two groups, Xu and colleagues and Yang and colleagues, demonstrated synaptic modification *in vivo* after mice learned a new motor behaviour, either a reaching task [109] or after rotarod running [110] (Figure 1.14). Although the murine motor system is less well-developed compared to the primate motor system, mice are still adept learners. Indeed, in the task used by Xu and colleagues, mice learned to reach around an obstacle for a food pellet with the forepaw, which is a substantial deviation from any behaviour performed in nature

and requires engagement of the rodent motor cortex, unlike most natural motor behaviours mice perform [109].

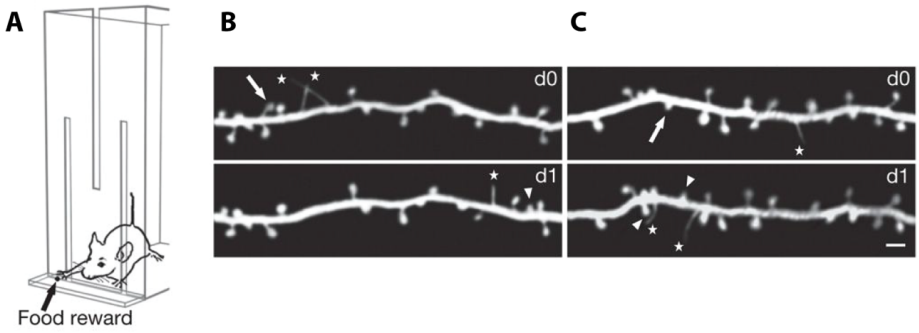


Figure 1.14. Dendritic spine formation in mouse motor cortex after learning a novel reaching movement. A. In a study by Xu and colleagues, mice were taught to reach through a narrow window using the forepaw to obtain a food reward. Dendritic spine formation (arrow heads), elimination (arrows), and filopodia (asterisks) in a control mouse (B) and a trained mouse (C). Top row shows images from day 0, prior to training. Bottom row shows images after 1 day of training. Alteration in dendritic spines was significantly higher in trained mice, and level of alteration was correlated with the degree of success on the task (data not show). Figure adapted from [109].

1.3.3 Population tuning for learning

What purpose does synaptic modification serve? One possibility is that the added synapses change the tuning properties of individual neurons. A second possibility is that the tuning of individual neurons remains constant, and the synapses effectively change how the population of neurons is combined. Computational models of perceptual skill learning suggest that the latter is closer to the truth. Tuning of individual neurons is altered in the context of a population, does not produce learning. Instead, robust learning can be achieved when individual neuronal tuning properties remain constant and the weights between the neurons are allowed to change [111, 112]. Therefore, although synaptic modification may underlie plasticity processes, in most cases, any single change may be inconsequential. Thus, understanding how a new behaviour is represented

as learned requires seeing how the activity of the population of neurons has changed [113].

In order to perform different behaviours, unique sets of neurons in a population are activated (Figure 1.15). So, to perform one behaviour, such as sweeping, a specific set of neurons in the motor cortex become active. Performing a different behaviour, such as playing the guitar, requires activating a different subset of cells from the motor cortex. In this scheme, learning a new action requires a new population code.

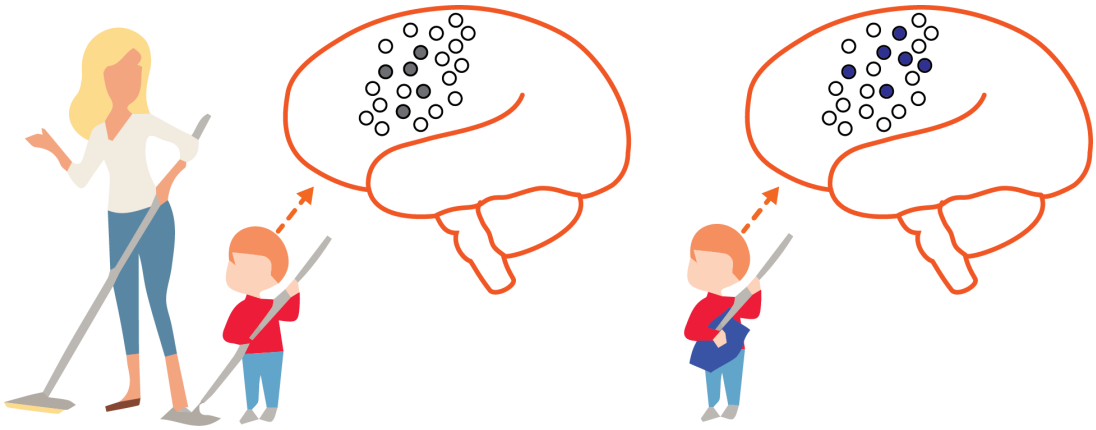


Figure 1.15. Performing an action requires activating a set of neurons in a population. When a child, Eddy, performs one action, for example sweeping, a set of neurons in motor cortex (circles) become active (grey filled circles). Later, a different subset of the population (blue filled circles) become active when Eddy plays the guitar.

How might a new population code be formed? One theory posits that a population of cells undergoes three distinct phases: 1) exploration of possible neural representations, 2) selection of neural representation that produced the most optimal outcome, 3) and refinement of the neural representation [33, 113]. The process of exploration, selection, and refinement produces an elegant symmetry between the variability evident in behavioural performance and the variability in the activity of neurons in a population (Figure 1.16). As an example,

behaviourally, if a child is learning to sweep, they might first explore many different ways to use the broom. Exploration would lead to a high-degree of variability in where the dust ends up. After exploring many different possibilities, the child will eventually converge on a way of holding the broom most effective at getting dust into the dustpan, thereby reducing the time spent sweeping. As the child repeatedly sweeps throughout their life, they will continue to refine their sweeping technique, reducing the variability in the amount of time spent sweeping, while the mean time spent sweeping remains constant. Thus, over the course of learning, variability goes from high to low [113].

Neuronally, the explore, select, refine hypothesis predicts a similar trajectory in variability. During exploration, the subset of neurons in a population active during behavioural performance increases in variability. Successful attempts lead to the selection of a final set of active neurons that become consistently active during performance. As the skill is refined, the recruited neural population also decreases in variability [33, 114]. In the next section, I provide additional detail of this account.

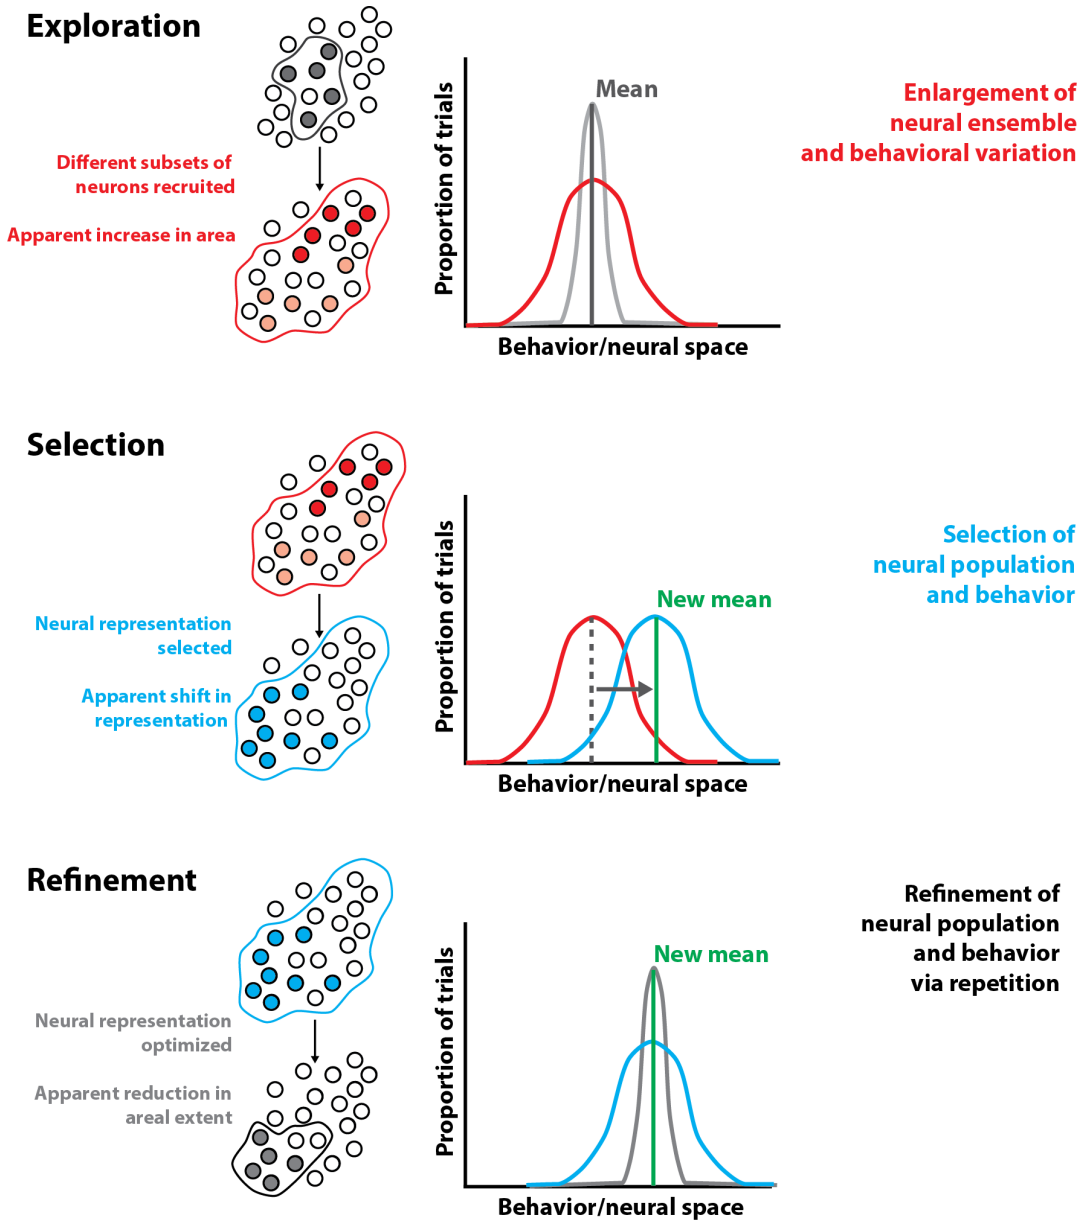


Figure 1.16. Stages of population tuning. During the exploration phase, behavioral performance is highly variable. Concurrent with the high behavioral variability, the neural ensembles activated are also increased (red bell curve). This creates an apparent increase in cortical area representing the behavior (red outline). In selection, the neural population that led to the most desirable outcome (blue) forms the new representation for the skill. This results in a shift in the behavior produced toward a new form, and the average neural population recruited during skill performance (blue bell curve). During refinement, the neurons activated from the selected population is gradually winnowed away, further reducing behavioral and neural-ensemble variability (grey bell curve). Figure redrawn from [113].

1.3.3.1 Exploration

Exploring subsets of neurons within a population increases the apparent cortical area devoted to the representation for a particular skill [33]. This apparent increase in area after initial learning has been observed in electrophysiological recordings in non-human primates [32, 115]. In these studies, the initial area representation of the forelimb in primary motor cortex was mapped using intracranial micro stimulation. Monkeys were subsequently trained on a motor task that required coordinated forelimb movements (bolt turning). When the representation of the forelimb was mapped after training, the researchers observed the cortical area that could evoke forelimb movements had increased in size. Similar increases in the apparent areal extent of the cortical representation of the forelimb after training have been observed in humans using transcranial magnetic stimulation [116, 117] and fMRI [8, 118].

The physiological mechanism driving the neural pattern exploration is not well understood. One potential mechanism to drive exploration is a reduction in local inhibition, which "unmasks" latent connections within a population [119] and may increase the variability in neural firing [113]. Reduction of the primary inhibitory neurotransmitter, gamma-aminobutyric acid (GABA), occurs during motor learning in humans [16, 120], and may result from alterations in dopaminergic signalling in the basal ganglia that change the balance between the excitatory and inhibitory output pathways [121-123].

1.3.3.2 Selection

Exploration allows more diverse sets of cells to 'fire' together [33, 109]. The subset of neurons that produced the optimal outcome can then be 'selected' as a final representation.

Selection requires input signals to relay the behavioural outcome. To this end, for motor skill learning, the basal ganglia may provide the prediction error and reinforcement learning signals to motor cortex, and the dopaminergic midbrain may also signal to motor cortex directly [124-127]. In addition, the cerebellum also provides an error-based teaching signal to motor cortex [128]. The representation in motor cortex is then stabilized by modulation of local inhibitory circuits [129].

1.3.3.3 Refinement

Following selection, the number of neurons representing the skill diminishes as the skill is refined [33, 113]. This reduction can lead to an apparent decrease of cortical area that represents the trained limb in motor cortex. The decreased number of neurons representing a well-learned behaviour has been observed in mice trained extensively on a forelimb reaching task and in monkeys trained on a visuomotor task [33, 130].

When examining the cortical area devoted to the representation of a skill using fMRI, a somewhat different picture emerges. Some studies have reported that long-term training causes increased BOLD-signal in primary motor cortex after long-term training, which may indicate that more neurons are activated during performance [8, 15, 131]. However, other studies have reported decreased BOLD-signal during performance after long-term training [132-134]. These

discrepancies may be due to differences in experimental designs across these studies, including the skill being learned or amount of training. In addition, the BOLD-signal can be affected by many other factors that pose challenges for interpreting 'activation-level' changes.

In addition, although the cortical area representing a particular movement may decrease as the representation is refined, developing long-term expertise may require cortical expansion to permit appropriate differentiation between multiple overlapping representations. Expansion of a cortical area has also been observed in animal models during the late phase of skill learning [135]. In addition, structural MRI evidence suggests this may occur in humans. For example, the hand area of expert pianists tends to be enlarged compared to amateurs [136, 137]. However, because studies of expertise in humans are often cross-sectional, causality is difficult infer from these experiments.

1.3.4 The bigger picture

Exploring, selecting, and refining population codes during the learning process explains many phenomena observed in experiments investigating skill learning. Yet some phenomena observed in population-based studies of plasticity are counter-intuitive. For example, if population tuning were required for learning, we might expect reshuffling of neural representations to occur with ease. However, experimental evidence suggests that, at least over short time-scales, neural representations in primates can be quite rigid. A seminal study by Sadtler and colleagues demonstrated this rigidity using brain-machine interface (BMI) learning in rhesus macaques [138]. In this experiment, monkeys learned to control a computer cursor by modulating neural activity recorded directly from their cortex using microelectrode arrays. By using this set-up, Sadtler and

colleagues could assess precisely which combinations of neurons in a population reliably co-activate. When the mapping of brain activity to the cursor followed pre-existing neural activity patterns, the animals were able to learn to control the cursor with ease. However, if the activity pattern was not among the pre-existing patterns, the animals learned less well or even failed to learn at all.

In a subsequent study, also using BMI learning in macaque, Golub and colleagues further demonstrated the rigidity of these patterns [139]. In this study, existing representations were stratified along two dimensions (Figure 1.17), such that half the neural representations would move the cursor to the left, and the other half moved the cursor to the right. After learning this mapping, the dimension that divided the leftward and rightward populations was then rotated to determine whether the underlying neural representations would be shifted to efficiently match the new mapping. In actuality, the authors found that the representations were not shifted; instead the representations remained constant and were re-associated to the new directions.

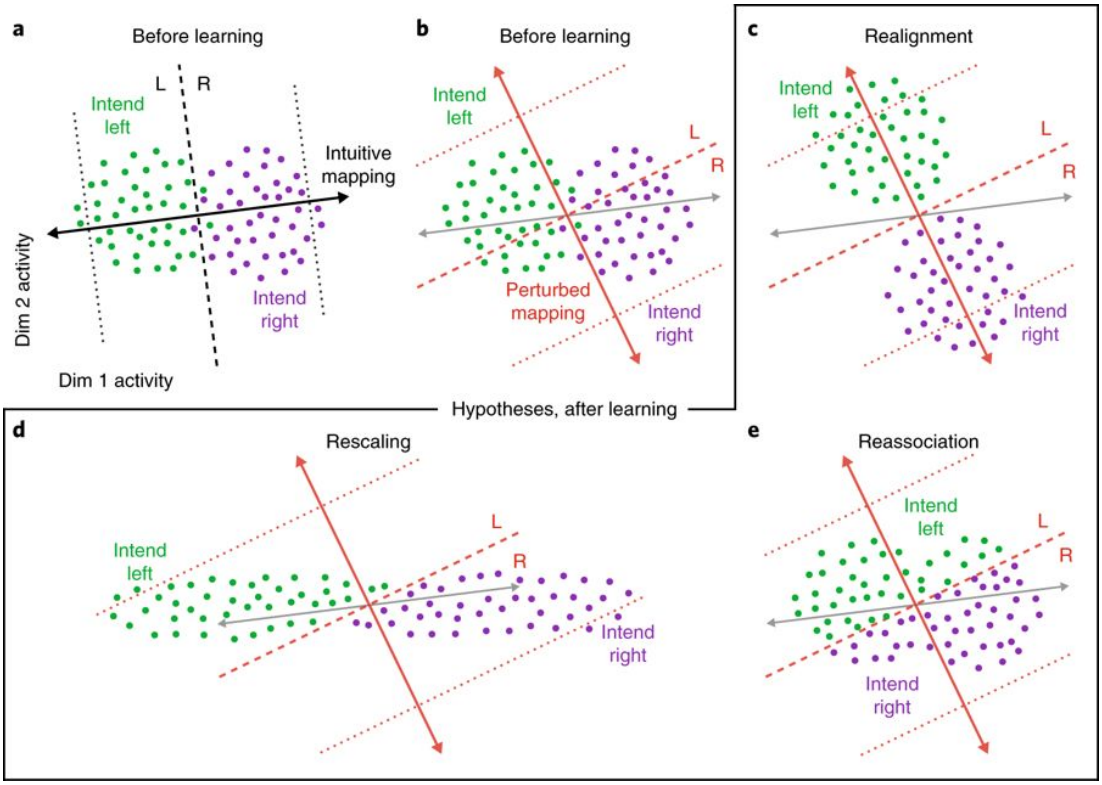


Figure 1.17. Learning by neural reassociation. This figure shows a conceptual representation of the findings from a brain-machine interface study in macaques by Golum and colleagues. A. Neural activity patterns are represented as pointes in a 2d 'neural activity space.' Initially, activity patterns to the left of the thick hashed line move the cursor left and are colored purple. Likewise, activity patterns to the right of the thick hashed line move the cursor right and are colored green. The distance from the thick dotted lines dictate the speed of travel speed. B. The dimension classifying the activity patterns was rotated (red lines), and monkeys needed to learn the new mapping in neural-activity space. C,D,E. Hypothetical activity pattern changes. C. The activity patterns would rotate to maximize behavioral performance. D. The rotation of the neural activity space causes the influence of dimension 1 to decrease and dimension 2 to increase. To compensate, the neural activity patterns scale along these dimensions such that the influence of dimension 1 increases and the influence of dimension 2 decreases. E. Activity patterns remain constant and are re-associated to the new direction. C and D predict that the neural representations would change, whereas E does not. The experimental data suggests E occurs in the primate brain. Figure reproduced from [139].

Finally, in a rodent model where new neural representations have been generated, one might expect that if the subset of neurons that will eventually come to represent the skill are activated at the beginning of learning, that the final form of the behaviour should be produced. However, this is not the case: there is very

little correspondence between behaviours produced by a given neural representation at the early and late stages of skill acquisition [33].

These oddities may be due to the fact that activity within a population of neurons is occurring in a larger network of interconnected populations. Each population may be separately tuned, as can the relationships between the populations. In such an arrangement, if representations within each population are being explored simultaneously, it may be difficult to predict the final representation based solely on the initial representation. Moreover, some populations may be more flexible than others. For example, a higher layer (e.g., PMC mapping onto primary motor cortex) may be more flexible than a layer that maps onto a more rigid representation (motor cortex mapping onto the spinal cord). These issues are explored more in depth in the following section.

1.3.5 Network-level activity changes

The studies by Sadtler and colleagues and Golum and colleagues suggest that, in the context of complex, hierarchically organized populations of neurons (layers), it may be more efficient to alter the interaction between the layers than to change representations within the layers. Electrophysiological recordings during visuomotor adaptation learning in rhesus macaques suggest that the motor system may operate this way [140]. Perich and colleagues suggest the motor system may adapt to visuomotor rotations in one of three ways: 1) change the activity of motor cortex neurons to remap the movement, 2) change the mapping of PMC neurons to motor cortex neurons (potent space), or 3) change activity among the neurons in PMC that alter movement preparation (null space; Figure 1.18). To test between these possibilities, Perich and colleagues trained monkeys to reach to targets, and recorded neuronal responses from the dorsal premotor

and primary motor cortices. The authors determined that the learning occurred in the null-space, suggesting that the representations remained constant but the manner those representations were initiated changed.

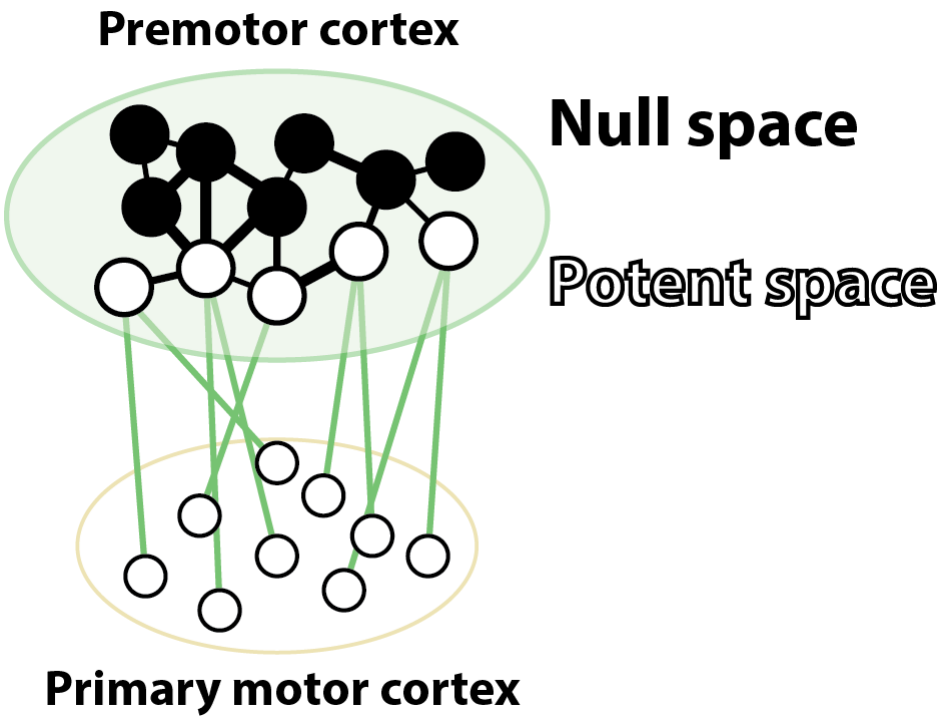


Figure 1.18. Diagrammatic representation showing the difference between the null space and potent space. The potent space refers to PMC neurons that activate specific sets of primary motor cortex neurons. The null space refers to neurons PMC neurons that do not directly excite primary motor cortex neurons. The study by Perich and colleagues [140] suggests that rapid visuomotor adaptation learning is facilitated by alteration in null space activity, rather than generation of novel activity patterns in the potent space of PMC, or the overall activity patterns generated by PMC or primary motor cortex.

It should be noted, however, that in the rodent brain, there is evidence for the generation of novel neural patterns. For example, Athalye and colleagues trained mice to generate completely novel neural activity patterns in a close-loop system similar to a BMI [114]. However, this set up was far from 'natural'; when the correct activity pattern was produced, the mouse received direct stimulation of the VTA. Therefore, while this study does prove that generating novel activity

patterns is possible, it requires a suitably long training period, a suitably narrow definition of success, and a suitably large reinforcement signal.

1.3.6 Summary

To summarize, structural and functional changes that facilitate learning are evident at multiple scales, from alteration of single synapses, to modifying local connections, to activity propagated between populations. While structural changes may occur as a product of long-term learning, the study by Perich and colleagues shows that short-term learning can occur by modifying activity in a network and does not require altering the structure of the network [140].

The most meaningful measurements of *functional change*, then, may be the interactions between brain regions. For this thesis, I define functional plasticity as *the alteration in the coordinated activity between two or more populations of neurons to promote a specific behavioural change*. In the next section, I explore studies that show modification of activity occurring across the brain during learning I conclude that conceptualizing the full dimensionality of the neural and behavioural space is therefore necessary to appreciate functionally-relevant and 'learning-related' changes in brain activity.

1.4 Functional plasticity

The study by Perich and colleagues [140] demonstrates that interactions between brain regions can promote learning – specifically, alteration in the pattern of signals sent from PMC to control primary motor cortex produced visuomotor adaptation. But, of course, the brain is composed of more than two layers. The question becomes: how does population remapping scale as layers are added? To address this question, Makino and colleagues [141] used whole-brain

calcium imaging in mice. Mice were trained to perform a forepaw lever pressing task, known to evoke population-level changes in M1 [33]. Whole-brain calcium recordings revealed that this task induced wide-spread changes in brain activity. Most notably, over the course of learning the whole-brain activity dynamics during the movement became more synchronized. This enhanced synchrony was primarily driven by M2, the mouse homolog to PMC [141].

These results suggest that learning produces wide-spread changes across the mouse brain. Several influential studies in humans have found similar functional plasticity-like changes in the human brain after motor skill learning [142, 143]. Using the co-fluctuation of the BOLD-signal as a proxy for neural communication (functional connectivity, see Chapter 2), graph-theoretical approaches suggest that the wide-spread activity changes promote information transfer between relevant brain regions during performance. In addition, the capacity of brain areas to be functionally connected with different regions (flexibility), was a strong predictor of participant's ability to learn [142]. In a subsequent study, the authors found that the flexibility of motor and visual brain areas was particularly important, which was also generally consistent with the results reported by Makino and colleagues [143].

In addition to studying functional plasticity during training, functional plasticity can also be assessed by examining changes in spontaneous brain activity before and after training (resting-state fMRI; e.g., [144]). This design avoids conflating BOLD-signal changes due to brain change with BOLD-signal changes due to performance changes that may be epiphenomenal, like time on task. For interpretational purposes, studies using resting-state fMRI assume that the change in correlated BOLD activity reflects, at least in part, ongoing neural

coupling during the post-training period (but see Section 2.2.3.2 for interpretational challenges in resting-state fMRI). In the next section, I discuss several important findings using this study design.

1.4.1 Functional plasticity after training

Whole-brain functional plasticity investigations conducted in the period after training have been largely consistent with the findings during training. Sami and Miall [145] found that the reorganization of the whole-brain functional connectivity network after training maximizes the local efficiency for information transfer. In a separate study, Taubert and colleagues showed that the SMA and ventral PMC facilitated this the network reorganization [146].

Several hypothesis-driven investigations have focused specifically on the functional connectivity of the cortical motor areas. Vahdat and colleagues found that increased functional connectivity among the cortical motor regions was positively associated with performance after short-term training [147]. In addition, two other studies focused specifically on motor network functional connectivity change using independent components analysis. In one study, functional connectivity of the motor network was assessed before and after participants trained for six weeks on a juggling task at either low-intensity (15-minutes per day) or high-intensity (30-minutes per day). After low-intensity training, functional connectivity within the motor network increased. In contrast, motor network connectivity decreased after high-intensity training [148]. Consistent with the result reported by Sampaio-Baptista and colleagues, Ma and colleagues reported that connectivity within the sensorimotor network increased after two-weeks of training but decreased after two additional weeks [149].

Finally, three studies have examined the impact of a single bout of training on functional connectivity. In all three studies, connectivity of a parietal-prefrontal network increased after motor learning [144, 146, 150]. None of these studies reported changes specifically within the motor network after short-term training.

Together, these studies show clear effects of motor skill learning on functional connectivity. The effects are fairly wide-spread. The functional connectivity of the motor network is reliably altered after multiple training sessions, while a single bout of training may also impact prefrontal cortex and parietal cortex functional connectivity.

1.5 Research questions

The research presented above has demonstrated that motor skill learning alters functional connectivity. In addition, the convergence between these results and invasive recordings in rodent models suggests that functional connectivity changes may serve as a marker of functional plasticity. However, at present, the literature investigating functional plasticity with fMRI has several limitations:

1. Most studies use only a single task. Therefore, it is not clear what aspects of these 'brain-wide' changes are more general, i.e., occur when learning any task is learned, and what changes are task-specific.
2. The extent to which laboratory tasks capture the effect of participant motivation may be limited. Others have found that external motivation (i.e., reward) is important during learning and can affect performance and memory [17, 151, 152]. Therefore, it is important to understand how these external influences are reflected in the brain during skill learning.

3. Finally, although there is evidence that regional inhibition may influence functional connectivity of brain regions, this has been restricted to measurements within a single region of interest (ROI) [153, 154]. Given the wide-spread effects of learning on functional connectivity, it is important to understand how the influence of local inhibition scales as more brain regions are considered.

1.6 Aims

My primary aim in this thesis is to probe the effect of short-term skill learning on brain activity using MRI in humans. I go about this goal by addressing the three main limitations described in Section 1.5. In this section, I outline the contents of the subsequent chapters.

In Chapter 2, I provide an overview of two important behavioural tasks used in this work. I also describe my resting-state fMRI and magnetic resonance spectroscopic imaging (MRSI) methodologies.

In Chapter 3, I address the first question posed above by investigating whether resting-state fMRI is sensitive to task-specific functional plasticity. In addition, I explore the specific interpretational challenges associated with studies of functional plasticity.

In Chapters 4 and 5, I address the second question posed above by investigating the influence of external motivation, i.e., rewards and punishments, on performance during motor skill learning. Then, in Chapter 6, I examine the effect of reward and punishment on functional plasticity in the brain, and the extent to which the influence of valenced feedback on brain activity is specific to task.

In the final experiment, presented in Chapter 7, I address the third question posed above by probing the association between functional connectivity and concentration of neurochemicals across the brain. For this study, I implement a novel MRSI technique that allows metabolite quantification across many brain areas simultaneously. I also begin to test the extent to which the relationship between the functional connectivity and GABA concentration can predict performance on a variety of sensorimotor tasks.

Finally, in Chapter 8, I summarize the findings presented in Chapters 3 through 7 and present directions for future research.

2 Methods

In Chapter 1, I examined the mechanisms that underlie learning and how changing interactions between brain regions can facilitate short-term learning. In this Chapter, I explore the methodology that I employed to study functional plasticity. I begin by discussing the behavioural tasks used. Subsequently, I discuss the principles of magnetic resonance for the purpose of magnetic resonance imaging. I then discuss the use of resting-state functional magnetic resonance imaging to measure training-induced changes in brain function. I conclude by discussing the novel magnetic resonance spectroscopic imaging technique used in the final Chapter.

2.1 Tasks for studying sequence learning

S*equence learning* refers to acquiring knowledge of temporal dependencies, and most motor skills require learning the order in which movements should be performed to accomplish the goal. In this thesis, I use two sequence learning tasks, the Serial Reaction Time Task (SRTT) and the Force Tracking Task (FTT). I describe these tasks in detail below. Practically, I chose these tasks in order to replicate previous work investigating the questions I raised in Chapter 1. In addition, these two tasks are complimentary because they require learning a sequence with distinct demands: execution speed (SRTT) versus precise control (FTT). In the SRTT, participants perform a movement sequence (generally using the fingers) as quickly as possible. In contrast, in the FTT participants perform a series of muscle contractions (again, generally with the hand or fingers) to follow a target trajectory. Using this dichotomy, I can investigate whether certain

principles (e.g., the impact of reward and punishment) generalize across different demands.

2.1.1 The serial reaction time task

The SRTT, conceived by Nissen and Bullemer (1987), has been among the most influential learning tasks in neuroscience and psychology [155, 156]. The SRTT features discrete cued movements (trials; Figure 2.1) that form a sequence. In the prototypical implementation of the SRTT, a participant is seated in front of a computer screen holding a response device (e.g., a button box). A set of locations (e.g., white ovals) are shown on computer screen, and the appearance of a stimulus (e.g., a white cross) in one of the locations cues the participant to press the corresponding response button.

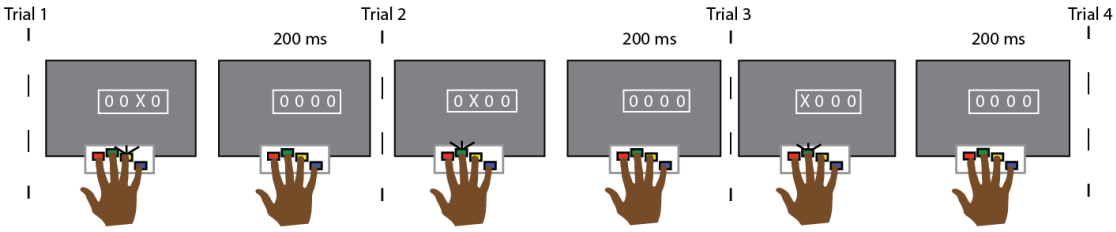


Figure 2.1. A prototypical implementation of the Serial Reaction Time Task. In the SRTT, a participant views a computer screen while holding a response device. A set of spatial locations (circles) appears on screen; a trial begins when one location changes (e.g., to a cross), which cues the participant to make the corresponding response. After the response, the cue is removed, and, following a delay (e.g., 200 ms), the next trial begins.

Knowledge for the sequence (sequence-specific knowledge) can be assessed by comparing the difference in RT for fixed-sequence versus random-sequence trials. Sequence-specific knowledge can be contrasted with general performance improvements, which are evaluated by comparing the RT random-sequence blocks at the beginning and end of training. In principle, both sequence-specific knowledge and general performance improvements can be assessed using

accuracy (number of errors) as an outcome measure. However, in practice, the accuracy on the SRTT is quite high, and thus RT is considered to be the most robust outcome measure [157].

2.1.1.1 A note on the original purpose of the serial reaction time task

Although it has been used to describe motor learning, the original purpose of the SRTT to provide its insight into the role of awareness in learning. While it was known that participants with anterograde amnesia could learn and retain a variety of simple motor skills, it was debated whether *explicit memory* was required for learning to take place [158]. Nissen and Bullemer sought to show that learning could take place in the absence of explicit memory or attention to the sequence [155, 156]. Using the SRTT, Nissen and Bullemer demonstrated that patients with Korsakoff's syndrome, which causes anterograde amnesia, could gain sequence-knowledge, despite having no ability to verbally recount the sequence they learned. This insight added support to the view that different mechanisms support explicit (knowledge that can be verbalized) and implicit (knowledge for skills) memory.

2.1.2 The force tracking task

To complement the discrete movements structure of the SRTT, I also implemented the FTT, which features a continuous sequence [15-17]. For the prototypical FTT participants sit at a computer screen with a force-transducing implement (generally held in the hand; Figure 2.2). At the beginning of a trial, a cursor (e.g., a white circle) and a target (e.g., a black circle) appear on screen. The target moves in one dimension (e.g., vertically), and the participant must modulate

their force to keep the cursor as close to the centre of the target as possible. The target continues to move until the trial is completed.

Tracking error is used to assess performance [15-17]. In my study, I used the squared distance from the center of the cursor to the center of the target. This enforced additional penalty for participants that strayed far from the target.

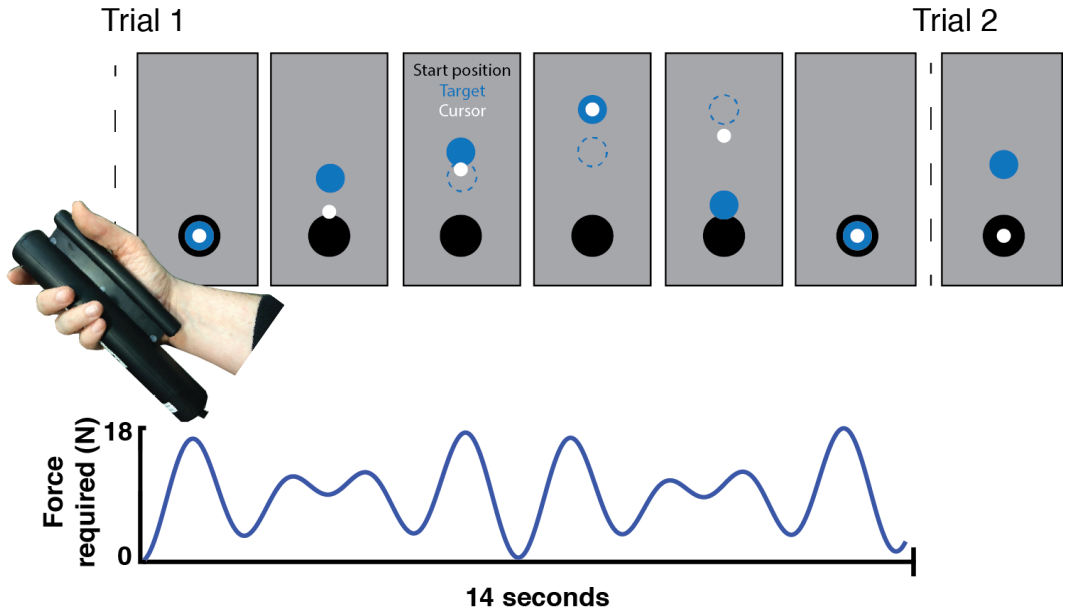


Figure 2.2. Trial structure in the Force Tracking Task. In the FTT, a participant sits at a computer screen holding a force transducing device. A target (blue circle) and cursor (white circle) appear on screen. The target moves vertically, and the participant must modulate their grip force to stay as close to the center of the target as possible. (Lower) During a trial, the target moves continuously on the screen; an example of a sequence of force application is shown in blue.

In the canonical implementation, the target follows the same trajectory on every trial [15-17]. However, this design may conflate general performance improvement with sequence-specific knowledge. To disambiguate these features and make this task qualitatively more similar to the SRTT, I implemented random-sequence trials in addition to the fixed-sequence trials. During the random-sequence trials, the target moved in a trajectory that was matched for amplitude and complexity with the fixed-sequence. Consistent with the SRTT, in this design,

I can assess sequence-knowledge by comparing the difference between the median squared-error during fixed-sequence trials versus random-sequence trials, and I assessed general performance by comparing performance at the beginning versus the end of training.

2.1.3 Comparison across tasks

In Chapters 4 through 6, I investigate the impact of reward and punishment on behavioural performance and neural activity evoked by training on the SRTT and FTT. This work was primarily motivated by two studies that reported a benefit of reward to memory using these two tasks [17, 152]. Although these tasks require dramatically different demands, they are often cited together to demonstrate the general principle that reward, and not punishment, benefits skill memory [151, 159-161].

Therefore, behaviorally, I wanted to replicate these studies. However, given the vastly different performance metrics between the two tasks, a direct, quantitative comparison of the outcomes was not practical. Although non-parametric statistics (e.g., Wilcoxen rank-sum tests) could be used, such methods would only allow me to compare the participant's performance relative to the group – effectively obscuring the outcome (benefit or detriment of feedback to learning). Therefore, in Chapter 4 (and the fMRI investigation in Chapter 6), I compare qualitatively across tasks, and make claims limited to the qualitative differences in the effect of feedback.

2.2 Measuring neural function in humans

In 1977, Damadian and colleagues demonstrated that magnetic resonance imaging (MRI) could be used to differentiate tissue types in the human brain

without administration of any contrast agents [162]. Since that discovery, MRI has afforded substantial insights into brain structure and function. The following section describes the principles of magnetic resonance (MR) physics and acquisition. Subsequently, I describe the preprocessing of MRI data relevant to my work in this thesis.

2.2.1 Concepts in magnetic resonance

Unlike X-rays or computed tomography that require delivery of ionising radiation, MRI leverages quantum properties of atoms in a sample to derive a signal. Hydrogen atoms (^{1}H, also referred to as protons) are most commonly used as the basis for the MR signal in neuroscience because of the high concentration of ^{1}H in biological tissue. Thus, while other molecules, such as phosphorous (^{31}P) and carbon (^{13}C), can be used for MRI, for simplicity I will refer specifically to proton MR.

MR relies on the quantum phenomena *charge* and *spin*. Charge refers to the positive and negative magnetic poles (Figure 2.3) which form an axis (*magnetic moment)*. Spin refers to the proton's apparent angular momentum about the magnetic moment.

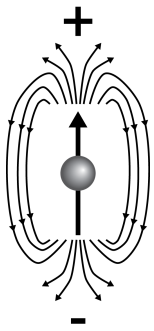


Figure 2.3. The magnetic moment. Dipolar molecules (e.g., hydrogen) exhibit positively and negatively charged poles. The direction of the positive-negative charge axis is referred to as the magnetic moment.

Typically, the protons' magnetic moments are randomly aligned; they interact with each other, but the sum of the magnetic moments (net-magnetization; M_0) is zero (Figure 2.4). When exposed to an external static magnetic field (B_0), the magnetic moments become aligned in parallel (low-energy) or anti-parallel (high-energy) with the external field. In this state, M_0 is equal to the difference between the proton energy states, which depends magnetic field strength and the exposure duration.

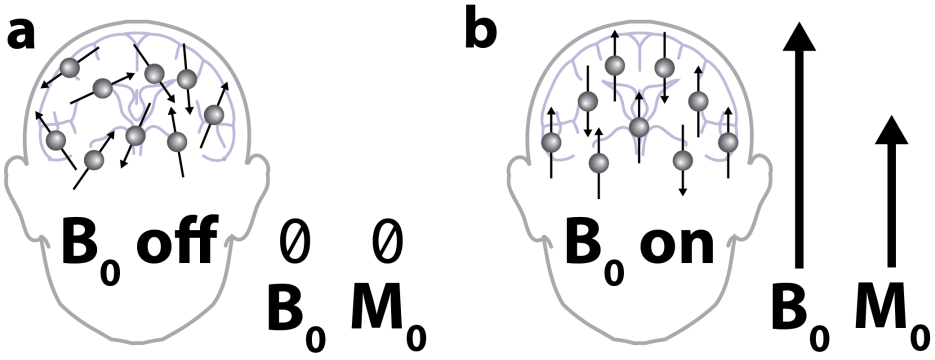


Figure 2.4. Magnetic dipole moments exposed to a strong magnetic field. A. Under normal conditions, the magnetic dipole moments in a sample are randomly aligned. B. When exposed to a strong external magnetic field (B_0), dipole moments align parallel or anti-parallel to B_0, creating a net magnetization (M_0), with the strength equal to the difference between parallel and anti-parallel dipole moments.

While aligned to the magnetic field, the protons spin (precess) at a specific frequency (Larmor frequency, $\boldsymbol{\omega}$) that depends on the magnetic field strength and the gyromagnetic ratio of the atomic nuclei. If exposed to radio frequency (RF) energy with a frequency that matches $\boldsymbol{\omega}$, a proton will absorb this energy and is 'flipped' out of alignment with B_0 (Figure 2.5). When the RF energy is removed, the proton will return to alignment with B_0 (relaxation), emitting RF energy as it does. The emitted RF energy of the relaxing protons induces current in conductive wire (receive coil) surrounding the sample, providing the basis for the MR image.

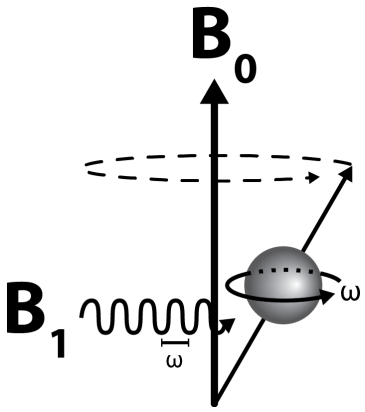


Figure 2.5. Precession and the Larmor Frequency. Protons exposed to B_0 spin (precess) at a specific frequency, the Larmor Frequency, based on their gyromagnetic ratio and the strength of B_0. These protons can absorb radio frequency energy at the Larmor frequency. This energy flips the proton's rotational axis away from B_0.

2.2.2 Magnetic resonance for imaging

In practice, the MRI scanner consists of a large magnet to generate B_0 to align the protons of the sample, yielding M_0 along a single axis (z-axis by convention; Figure 2.6). A second, smaller, oscillating magnetic field (B_1) is used to generate an RF pulse at the Larmor frequency. The RF pulse flips M_0 into the transverse (x-y) plane, increasing magnetization in the x-/y-plane and decreasing magnetization in the z-plane. When B_1 field is removed, M_0 relaxes away from the transverse plane, yielding an increase in magnetization along the z-axis, and a reduction in magnetization in the x-y plane, detected in the receive coil. During relaxation, these signals are referred to as T1 and T2, respectively. If allowed to fully decay, the observed signal in the transverse plane undergoes an exponential dampened oscillation (*free induction decay* [FID]), where the frequency of the oscillation is approximately equal to the Larmor Frequency.

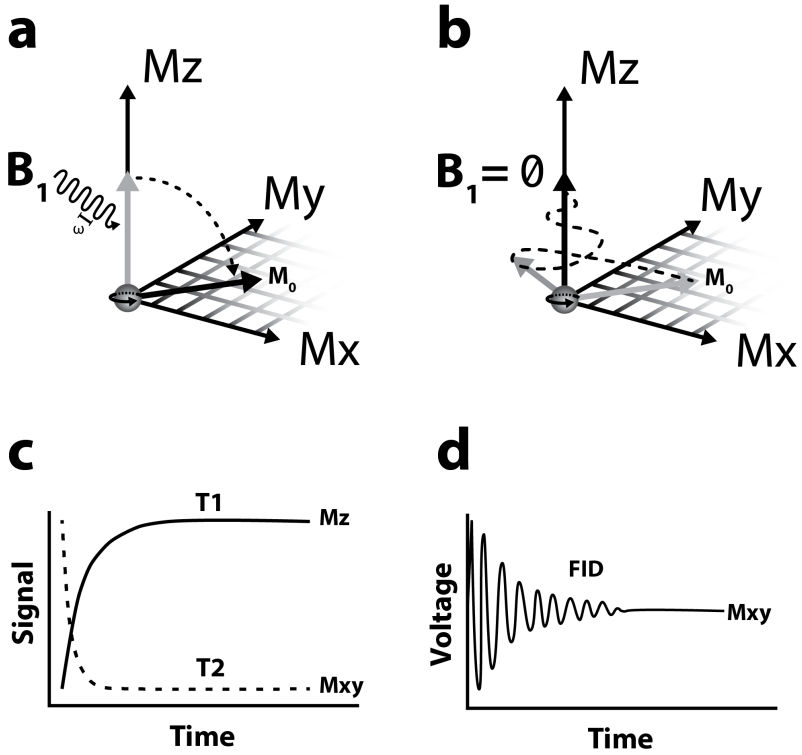


Figure 2.6. Nuclear magnetic resonance. A. Energy imparted by a radio frequency pulse (oscillating magnetic field B1) flips the net magnetization (M0) of the sample away from B0 (Mz, by convention), into the transverse (Mxy) plane. B. After B1 is turned off, M0 relaxes away from the transverse plane and grows in strength along the B0 axis. This change in magnetic field induces current in conductive material (receive coil), which can be quantified C. The return of M0 to B0 and relaxation away from the transverse plane are referred to as T1 and T2, respectively. D. The signal induced in the receive coil is a dampened oscillation with a frequency approximately equal to the Larmor Frequency; full measurement of this signal is the free induction decay.

2.2.2.1 Image contrast

The decay rate of the T1 and T2 signals depends on the content of the sample. These signals provide independent information about the content of the volumetric pixels (voxels) in the image. Contrast between tissue-types can be manipulated by varying the timing of the signal measurement after excitation (echo time, TE) and the time between excitation pulses (TR).

The rate of T2-signal decay is faster than the physics of MR would predict, due to local magnetic inhomogeneities, for example blood oxygenation level, that cause the protons to become imperfectly aligned during relaxation (dephasing).

This imperfect T2 decay is referred to as T2*. Measuring BOLD T2 signal changes is the basis for the majority of functional MRI (fMRI).

2.2.2.2 Image formation

Spatial localization is accomplished by imparting a spatially-dependent variation of the B_0 field using scanner's gradient coils (G_x, G_y, and G_z), which creates a spatially-dependent variation of precession frequency. Because of this spatial variation, spatially-specific excitation can be achieved with the matching RF pulse frequency (slice-selection).

In order to encode additional spatial information for the image, the other two gradients (G_x and G_y) are applied to encode position information in the high-dimensional Fourier-space, or k-space. The gradients impart spatial information in both the phase and frequency of the signal as its magnitude is measured. Following acquisition, a series of inverse Fourier-transformations are used to reconstruct the image from the measured signal. Different methods for sampling k-space can minimize acquisition duration, increase signal-to-noise, or improve fidelity of spatial localization.

2.2.3 Functional magnetic resonance imaging

Rapid MRI acquisition of T2* changes induced by the delivery of oxygenated blood to 'active' neural tissue can serve as a proxy for neural function.

2.2.3.1 The blood oxygenation level dependent signal

The BOLD-signal arises because oxygenated and deoxygenated haemoglobin have opposite magnetic properties: oxygenated haemoglobin of is diamagnetic, whereas deoxygenated haemoglobin is paramagnetic [163]. Active neural tissue consumes more oxygen than less active tissue. The increased activity

creates a local increase in the relative deoxygenated:oxygentated haemoglobin ratio. In response to decrease in oxygenation, the local capillary beds dilate, causing a rapid influx of blood with oxygenated haemoglobin. The increase in oxygenated haemoglobin vastly exceeds the amount required by the tissue. Following the decrease in neural activity, the local blood oxygenation remains at an elevated level for some period of time (approximately 8-12 seconds). This gives a hemodynamic response its stereotyped shape (Figure 2.7): an initial dip, followed by a large, prolonged increase in signal [164].

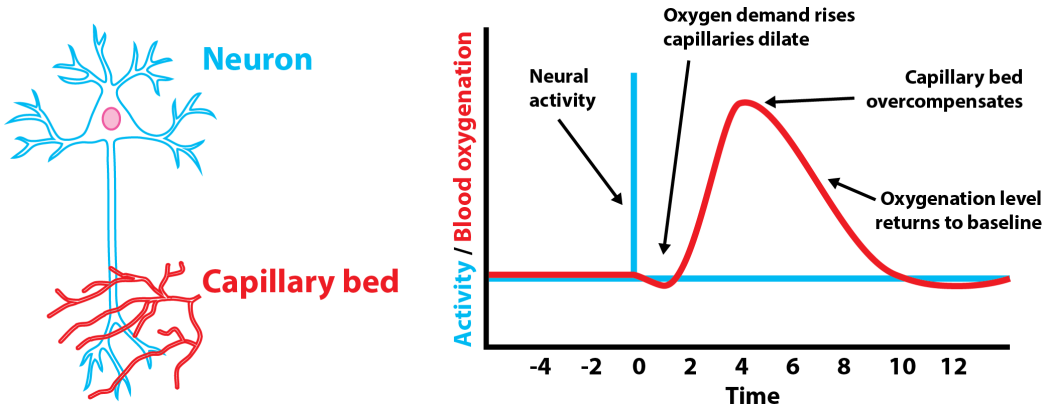


Figure 2.7. Time course of the hemodynamic response. Neural activity creates a local increase in deoxyhaemoglobin lasting approximately 1-2 seconds. In response, capillary beds dilate, bringing in oxyhaemoglobin that overcompensates for local demand, peaking between 4-6 seconds after the activity onset. Blood oxygenation level gradually returns to baseline, with a small undershoot. This process takes roughly 12 seconds, depending on local vasculature.

Seiji Ogawa is generally credited with discovering BOLD-contrast MRI [165]. When combined with rapid k-space sampling techniques, such as single-shot echo-planar imaging (EPI), T2*-contrast could be used to image neural function [166, 167]. In 1992, three research articles demonstrated the viability of BOLD-contrast fMRI for measurement of brain activity in humans [168-170]. Both Kwong and colleagues and Ogawa and colleagues demonstrated the increase in T2*-signal in visual cortex during visual stimulation, while Bandetinni and

colleagues demonstrated that hand movements increased T2*-signal in motor cortex (Figure 2.8). Although the BOLD-signal does not have the temporal resolution of electrophysiology, the ability to perform non-invasive whole-brain imaging opened up a new frontier for neuroscience, and many cognitive functions have since been localized to discrete regions of cortex using this technique (e.g., the fusiform face area, [171]).

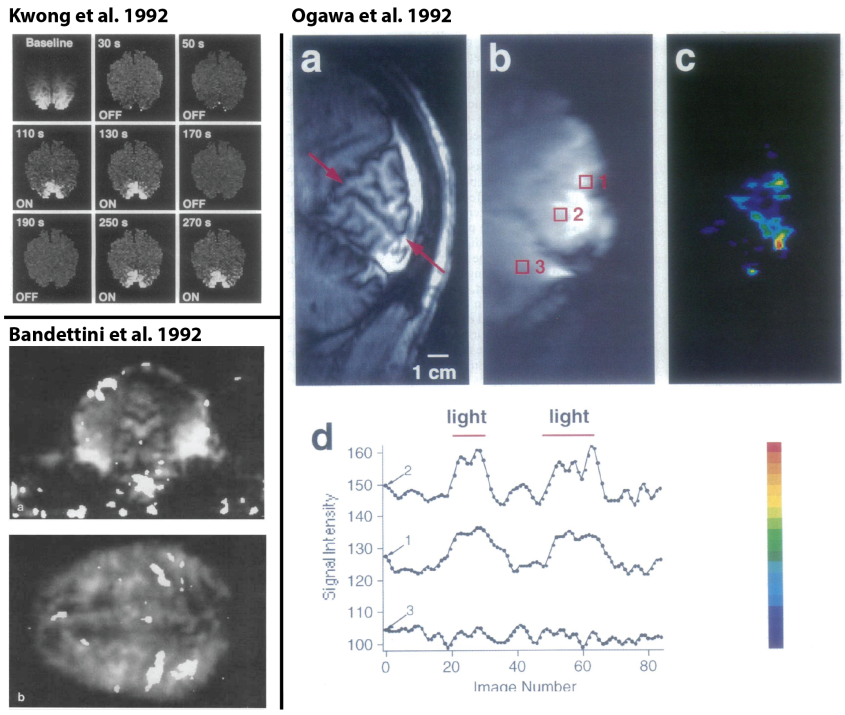


Figure 2.8. Demonstrating the feasibility of BOLD in humans. Both Kwong and colleagues and Ogawa and colleagues demonstrated the BOLD response of the visual cortex during visual stimulation. Bandettini and colleagues demonstrated the BOLD response of motor cortex evoked by movement. Figures adapted from [168-170].

2.2.3.2 Spontaneous BOLD-fluctuations, or 'The Resting State'

The BOLD-signal shows fluctuation even without stimulation or task-performance. In 1995, Biswal and colleagues reported that these spontaneous fluctuations were highest in functionally-related regions (those known to reliably 'coactivate' during task performance, e.g., bilateral motor cortices [172]). The correlated BOLD-signal during rest has been used elucidate functional

interactions in the healthy brain [173], to examine functional consequences of learning [150, 174], and to investigate changes in functional interactions associated with disease [175, 176]. Correlated BOLD-signal is often referred to as 'functional connectivity,' and I adopt this convention here.

Understanding the physiological origin of spontaneous BOLD-signal fluctuations is of great interest, and the extent to which neuronal activity drives spontaneous BOLD-signal fluctuations has been challenged. For example, one study [177] using a rodent model found that cerebral blood volume changes (CBV, measured optically) in somatosensory cortex was correlated with neural activity (measured by electrophysiological recordings) during volitional whisker and body movements. However, CBV change was only weakly related to neural activity during 'true' rest periods. In addition, spontaneous CBV changes were observed even when neural activity was pharmacologically inhibited, leading the authors to assert that neural activity did not drive spontaneous CBV fluctuations [177].

However, other studies strongly support a neural origin of the spontaneous fluctuations. One particularly important study by Ma and colleagues investigated the relationship between blood flow and neural activity using a transgenic mouse line that enabled differentiation of oxygenated and deoxygenated haemoglobin during optical recordings [178]. In this study, Ma and colleagues recorded the spontaneous fluctuation of oxygenated and deoxygenated haemoglobin while simultaneously imaging neuronal activity using a calcium reporter. The authors reported that, even when the animal was resting, the time course of regional oxygenated:deoxygenated haemoglobin ratio closely matched the time course of neural activity (Figure 2.9). The correspondence appeared to be driven largely by

the influx in oxygenated haemoglobin rather than any change in deoxygenated haemoglobin [178].

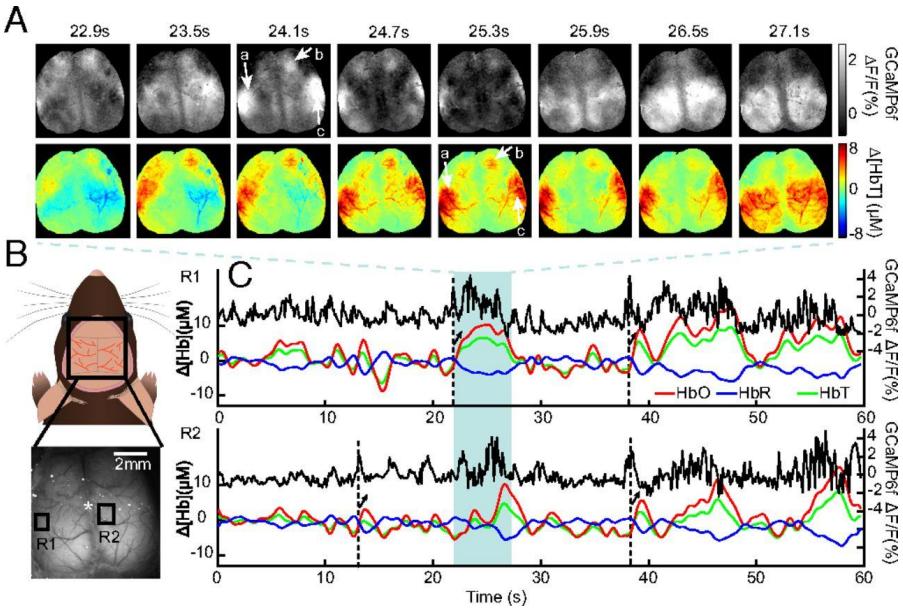


Figure 2.9. Blood oxygenation level closely matches spontaneous neural activity in rodents. A. Consecutive maps of fluorescent reporter of neural activity (GCaMP6f; top) and associated change in blood oxygenation (ΔHbT) viewed through thinned cranial bone window. Areas of increased neural activity (e.g., arrows a-c) show concurrent increases in blood oxygenation level delayed by 1-2 seconds. B. Example of mouse preparation. C. Time courses of GCaMP6f and blood oxygenation level change from two regions. Dotted lines indicate neural events, arrows indicate blood oxygenation level change. Black: GCaMP6f, Red: Oxyhaemoglobin, Blue: Deoxyhaemoglobin, Green: Total blood oxygenation level. Figure reproduced from [178].

Concurrent electrophysiological and fMRI data recorded from anesthetized monkeys provide further evidence for a neural origin of resting-state BOLD-signal fluctuations [179]. Murayama and colleagues found that spontaneous electrophysiological activity in primary visual cortex was related to the hemodynamic response with a 5s lag. In addition, this relationship was strongest in gamma and high-gamma frequency bands, which are thought to reflect local neural processing [180-182].

All available evidence suggests that spontaneous fluctuations probably arise from both neuronal and non-neuronal sources, and differentiation between oxygenated and deoxygenated haemoglobin may be key to determining the neural component [178]. Therefore, functional connectivity analyses critically depend on the estimation of T2*, as fluctuations about the mean can serve as a proxy for oxygenated:deoxygenated haemoglobin ratio. Thus, removing other noise sources must be achieved through effective preprocessing.

2.2.3.3 Sources of noise in resting-state fMRI: Motion

Motion is perhaps the most notorious source of noise in resting-state fMRI data [183-185]. Because of the nature of MRI acquisition, motion causes several issues that can be mitigated to varying degrees. The most obvious problem caused motion is misalignment of images. Volume registration during preprocessing helps to address this issue. Slice-based sampling, as is most common, causes another motion-related issue: sharp movements during an acquired volume will cause different areas of tissue to be over- or under- sampling of some regions; these artefacts are best corrected for by removing (censoring) affected time points [186].

A somewhat less intuitive motion-related problem arises from the physics of MR. MRI assumes the precise localization of the sample in the B0 field to enable spatially-selective slice localization. Participant movement violates this assumption by 1) changing the location of protons that precessing at excited frequency, 2) distorting the B0 magnetic field, and 3) causing dephasing of proton spins. These issues, collectively referred to as spin-history artefacts, cause immediate signal changes, which can last for multiple TRs (Figure 2.10).

During acquisition, the B0 field distortion and change in participant location can be mitigated by prospective motion correction during acquisition (e.g., PROMO; [187]). During preprocessing, strategies to correct spin-history artefacts at the participant level include correcting outlying time-points (despiking) or removing timepoints affected by motion.

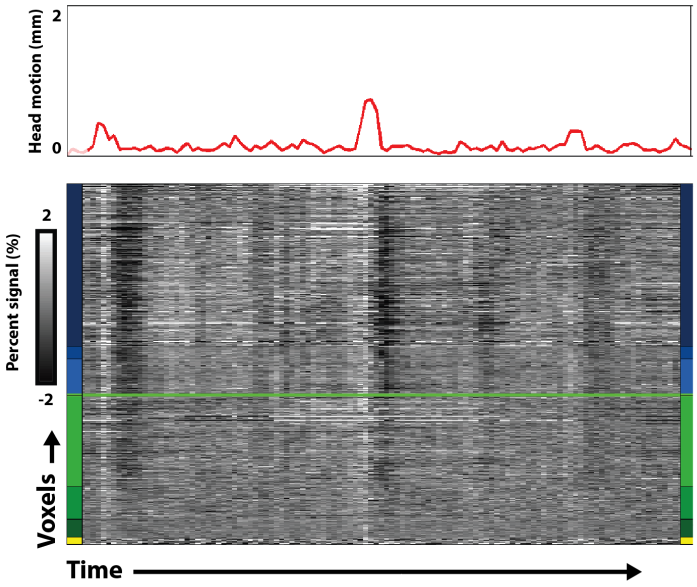


Figure 2.10. Motion induced global artefacts in fMRI data. Framewise displacement, a measure of total head motion is shown in the upper panel. Below is the time course from all voxels. Instances of large motion are followed by periods of widespread signal decreases lasting multiple TRs. These widespread signal changes can confound correlation analyses. Adapted from [183].

2.2.3.4 Sources of noise in resting-state fMRI: physiological

Physiological artefacts arising from the cardiac and respiratory cycles interact with the BOLD signal [188, 189]. These cycles cause pulsatile motion of the brain and T2* fluctuations because of the variable CO_2 concentration across these physiological cycles (which changes blood oxygenation level).

The TR determines how well these artefacts may be corrected. Imaging at a TR slower than the cardiac cycle, e.g., 2 seconds, causes this signal to be aliased into frequencies around 0.1 and 0.2 Hz, and a low-pass filter can remove this signal

[190]. However, more elegant solutions can be applied if physiological recordings are collected. The influence of differential CO_2 can be modelled and regressed out of the signal at the volume (respiration volume over time, RVT [188]) or slice level (RETROICOR ,[191]). In addition, because of their stereotypical frequency and spatial signatures, anatomy-based regressors (ANATICOR, [192, 193]) and ICA denoising techniques (ICA-FIX: [194]; me-ICA: [195]) can also remove physiological noise.

2.2.3.5 Sources of noise in resting-state fMRI: hardware

Hardware-related artefacts include low-frequency signal drift, which is generally removed using a high-pass filter. In addition, EPI-specific artefacts, including distortion and dropout, can corrupt data. Correction for EPI-based artefacts during preprocessing can be accomplished by 'unwarping' the EPIs based on a B_0 field map.

Acquiring data with using a short TE can also mitigate dropout. Multi-echo EPI, whereby multiple full volumes are acquired after a single excitation, can also be used to compensate for different optimal TE across the acquired image [195, 196]. Single-echo EPI uses a single TE that optimized for the contrast between de-/oxyhaemoglobin across the whole brain (Figure 2.11). However, areas with susceptibility artefacts caused by high-iron content (e.g., the striatum) or proximity to air-tissue interfaces (e.g., orbitofrontal cortex) generally have shorter optimal TEs and areas with signal hyperintensity due to coil proximity (e.g., visual cortex) may have a longer optimal TEs. By acquiring multiple images during the T2* signal decay from a single excitation, multi-echo EPI enables calculation of an optimal TE for each voxel based on a logarithmic fit across the

TEs. Based on this fit, an optimal image can be approximated. This substantially improves signal-to-noise (SNR) in areas with dropout.

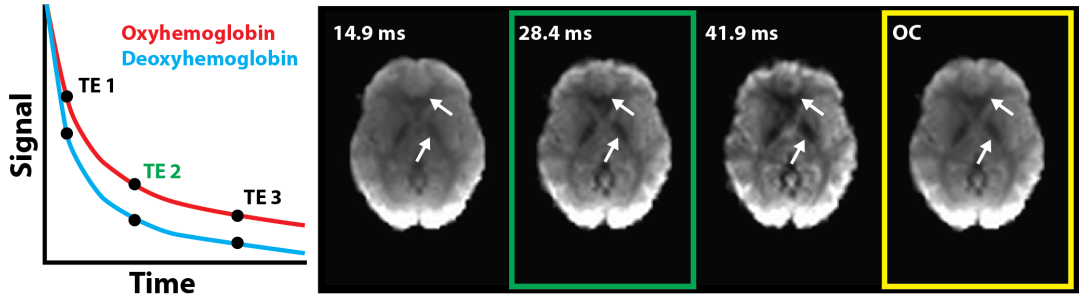


Figure 2.11. Multi-echo EPI improves SNR. Left. The T2-signal decays exponentially after excitation. For single-echo EPI, data from a single TE, optimized for the contrast between oxyhaemoglobin (Red) and deoxyhaemoglobin (Blue) is collected. This TE is generally around 30 ms (green). Right. Variation in signal intensity in EPIs acquired at multiple TEs from a single excitation. Regional variation in magnetic susceptibility, e.g., iron deposition or proximity to air cavities (arrows), causes faster T2 decay and signal-drop out in these regions at longer TEs. Signal hypo-intensity can pose problems for quantification in these regions at the average optimal TE (Green). Acquisition of multiple echoes enables calculation of the optimal TE at each voxel (Yellow).*

In addition to boosting SNR, multi-echo EPI can be used for the differentiating BOLD from non-BOLD signal during preprocessing [195, 197]. In Chapter 6, I use multi-echo EPI to boost image SNR. However, as the differentiation of BOLD from non-BOLD is an area of ongoing research, I did not use the multi-echo based denoising in this thesis.

2.2.4 FMRI preprocessing: A tale of two pipelines

In this section, I present several preprocessing choices relevant to the present thesis (shown in Table 2.1.) I used both the Analysis of Functional NeuroImages (AFNI, [198]) and FMRIB Software Library (FSL, [199]) packages.

Table 2.1. Preprocessing pipelines used for fMRI studies. Rows ordered by the sequence of steps applied

AFNI 1 (Chapter 3)	AFNI 2 (Chapter 6)	FSL (Chapter 7)
		B_0 unwarping
RETROICOR		
Time-shift correction	Time-shift correction	
RVT regressor calculation	RVT regressor calculation	
Despiking	Despiking	
Motion correction	Motion correction	Motion correction
	Echo combination	
ANATICOR	ANATICOR	
Smoothing	Smoothing	
Bandpass filtering	Bandpass filtering	
Nuisance projection/censoring	Nuisance projection/censoring	
		Melodic ICA decomposition
		FIX-ICA denoising

2.2.4.1 AFNI preprocessing

Physiological noise was corrected using RETROICOR (3dretroicor, AFNI) in the study in Chapter 4. Because of the complex interleaving of slices in the multi-echo acquisition, RETROICOR could not be applied in Chapter 6. In addition to RETROICOR, in both Chapters 4 and 6, I modelled the differential CO_2 concentration due to heart-rate and respiration using RVT (RetroTS, AFNI). Because RVT models CO2 concentration at the image-level, it can be applied to multi-echo data.

Following physiological noise correction, I applied slice-time correction. Slice-time correction (3dTshift; AFNI) resolves a critical feature of slice-based acquisition: the whole volume is not acquired at the same time. Many EPI acquisition schemes acquire slices in an interleaved fashion (for example, odd slices acquired, then even slices) to avoid cross-slice excitation. However, this

causes adjacent slices to be acquired at different time points in the signal decay. To compensate, 3dTshift interpolates the value across slices. For acquisitions with long TRs (e.g., 2 s; Chapters 4 and 6), slice-time correction improves consistency of analysis [200].

I next applied several steps to mitigate motion artefacts. First was *despiking*. Despiking (3dDespike, AFNI) reduces the magnitude of outlying time points. This process involves fitting a high-order polynomial to the timeseries data and evaluating the magnitude of the residuals from the fit. Time-points with outlying residual values are interpolated based on the model. This procedure improves the SNR of resting-state MRI scans and helps eliminate spurious distance-dependent correlations induced by motion [192].

After despiking, *motion correction* or *volume registration* (3dvolreg, AFNI) was performed. Here, volumes of the timeseries were registered to a reference time point (e.g., first volume). This registration is often done using six parameters (three rotations and three translations. The registration parameters for each time point are also regressed out of the data at the final stage of preprocessing.

Following motion correction, I applied the *ANATICOR* (@ANATICOR, AFNI) procedure. ANTICOR leverages the presence of noise in the white-matter and cerebrospinal fluid (CSF). In the white-matter, hardware-related noise dominates the signal from this tissue; taking the average signal from white-matter reveals high- and low-frequency drifts, and some spin-history related signal disruptions. The average white-matter signal can be used as a nuisance regressor. In addition to the average white-matter signal, the specific high- and low- frequency noise at a given grey matter voxel can be estimated based on the signal of the neighbouring white-matter, and regressed out of the signal [192, 193].

The CSF signal from the ventricles and sagittal sinus is influenced by pulsatile motion and respiration. ANATICOR correction also generates the average ventricular signal timeseries to remove these physiological noise sources via regression at the final preprocessing stage.

Finally, *temporal filtering* and *noise projection/deconvolution* was applied. For nuisance projection (3dTproject), all noise timeseries (motion parameters, physiological regressors, white-matter and CSF signals, and temporal-filter regressors) are projected out of the data, leaving the residuals as the cleaned timeseries. The cleaned data is then *spatially smoothed* to improve voxel-wise SNR. Smoothing also ensures that the data fulfil assumptions of Gaussianity for cluster-correction during group level analyses.

2.2.4.2 A note on echo combination

In Chapter 6 I acquired data using a multi-echo EPI sequence and optimally combined the echoes to boost SNR. The optimal combination procedure is described below [195, 197]:

The signal at each echo (n) varies as a function of the initial signal intensity S_0 and the transverse susceptibility T2* = $1/R_2$*, given by the monotonic exponential decay:

$$S(TE_n) = S_0 exp(-R_2 * TE_n)$$

Where R_2* is the inverse of the relaxation time, or 1/T2*. This equation can be linearized, simplifying the estimation of T2* and S_0 as the slope using log-linear transformation. The time courses can be optimally combined by summation, weighted by a factor, w, described by the following equation:

$$w(T^*_{2(fit)})n = \frac{TE_n * exp(-TE_n/T^*_{2(fit)})}{\Sigma n\, T^*_{2(fit)} TE_n * exp(-TE_n/T^*_{2(fit)})}$$

Where $T^*_{2(fit)}$ is the transverse relaxation time estimated for each voxel. The optimally combined timeseries is subsequently treated as a single echo.

In multi-echo preprocessing, all preprocessing steps prior to motion correction are applied to each echo independently. Motion correction is the computed based on the middle echo and applied to the other echoes. The preprocessing steps after motion correction are then applied to the optimally combined timeseries.

2.2.4.3 Preprocessing data with FSL

I used FSL to process the data in Chapter 7. I applied motion correct (mcFLIRT) and removed spatial distortion in the EPI based on a B_0 field map. I then registered the data to the high-resolution image (FNIRT). I then used FSL's Independent components analysis (ICA) to decompose the data into signals and sources, classified noise sources, and removed these from each participant's data [194, 201].

ICA is a dimensionality reduction technique that decomposes multivariate data such that independence in a particular dimension is maximized. In fMRI, ICA is generally used to decompose the spatio-temporal data into spatial sources of signal and their corresponding timeseries (MELODIC). These components generally have spatial and temporal properties that are characteristically BOLD-like or non-BOLD-like (Figure 2.12). For example, BOLD-like components are restricted to grey matter and have more power in the range 0.01-0.1 Hz. In contrast, non-BOLD-like components may be localized to white-matter or the ventricles and have greater power at high frequencies.

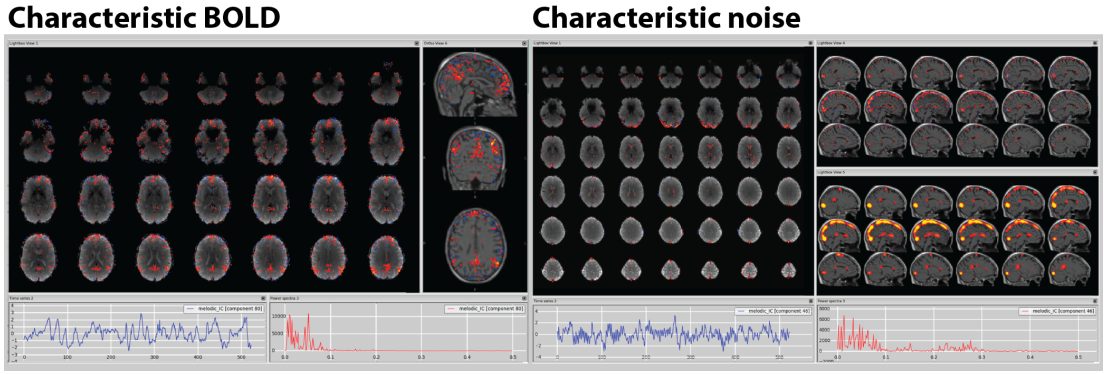


Figure 2.12. Characteristic BOLD-like (Left) and noise-like (Right) components. In general, BOLD-like components may be restricted to grey matter. In contrast, noise-like components have more often restricted to vasculature, ventricles, and white-matter. In addition, noise-like components may have greater power at higher frequencies (greater than 0.1 Hz), evident in the Fourier decomposition of the component timeseries. Figure adapted from [201].

I did not apply bandpass filtering in the study described in Chapter 7 because power in high frequencies is one criterion for noise-selection [201]

I also did not apply spatial smoothing. Because I used an anatomical parcellation for analyses, I wanted to avoid smoothing data across parcels. Further, because I did not use cluster-based statistical thresholding in this study, the group-level data comparisons did not require the assumption of a Gaussian noise distribution across space.

2.2.5 Resting-state group analysis

The work in this thesis employs two major resting-state analysis strategies, broadly classified as model-based and model-free.

A model-based approach is most useful for to test hypotheses about specific brain regions. The most common way of implementing this approach is a ROI analysis. I use ROI-based analyses in Chapters 3 and 6. In an ROI-analysis, the timeseries from a preselected ROI is calculated. This timeseries is then related to the rest of the brain using correlation or multiple regression.

In contrast, model-free analyses are most appropriate in exploratory studies or in studies examining more global phenomena, such as organizational principles of the brain. Examples of model-free functional connectivity analyses include graph theoretical approaches [202] and group-ICA[203]. I employ the group ICA method in Chapter 7.

I describe these methods in the relevant chapters of this thesis.

2.3 Magnetic resonance spectroscopic imaging

BOLD fMRI provides information about the activity and organization of the brain. However, blood oxygenation level is quite distinct from the electrochemical signals that underlie neural activity, making mechanistic inferences difficult to draw. Proton magnetic resonance spectroscopy (MRS) is a non-invasive technique that addresses this limitation and can complement fMRI.

MRS enables the measurement of neurochemicals, such as N-acetylaspartate (NAA), Choline (Cho), and Creatine (Cr), as well as the neurotransmitters glutamate (Glu) and gamma-aminobutyric acid (GABA), in the human brain. To date, MRS experiments have generally recorded from one or two voxels generally 8 cc in volume with an acquisition duration of 5-20 minutes. In this thesis, I implemented a novel MRSI sequence to acquire an entire slab of voxels (48 x 48 voxels) simultaneously. This sequence achieves a higher resolution (0.5 x 0.5 x 1.5 cm = 0.375 cc; approximately 21.3-fold improvement in nominal resolution), while simultaneously shortening acquisition time (9 minutes) at a comparable SNR, thus enabling robust metabolite mapping *in vivo* at 3-Tesla (Figure 2.13). These improvements brought MRSI toward clinical feasibility, as I demonstrated by mapping 2-hydroxyglutarate, an important biomarker present

in low concentrations, in a patient with a high-grade brain tumour (Figure 2.14; [204]).

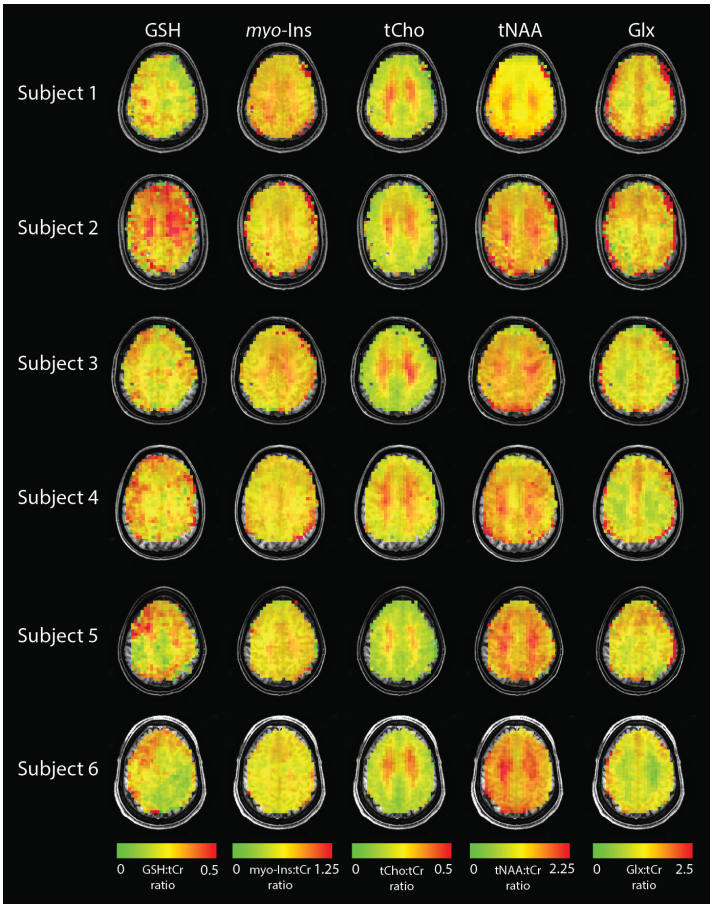


Figure 2.13. High-resolution MRSI acquired at 3-Tesla. Metabolite maps from six representative subjects acquired with the MRSI sequence used in the present thesis. The nominal resolution of 0.5 x 0.5 x 10 mm used here represents a 32-fold increase in resolution compared to current standard single voxel MRS acquisition. GSH: glutathione, myo-Ins: myo-inositol, tCho: total choline, NAA: N-acetylaspartate, Glx: glutamate and glutamine. All metabolites are presented relative to total creatine. Figure adapted from [204].

Given this technical advance, I sought to leverage MRSI to investigate the relationship between regional neurochemical concentrations and functional properties of the brain. In the following section, I cover basic information about MRS and details the MRSI sequence used in Chapter 7.

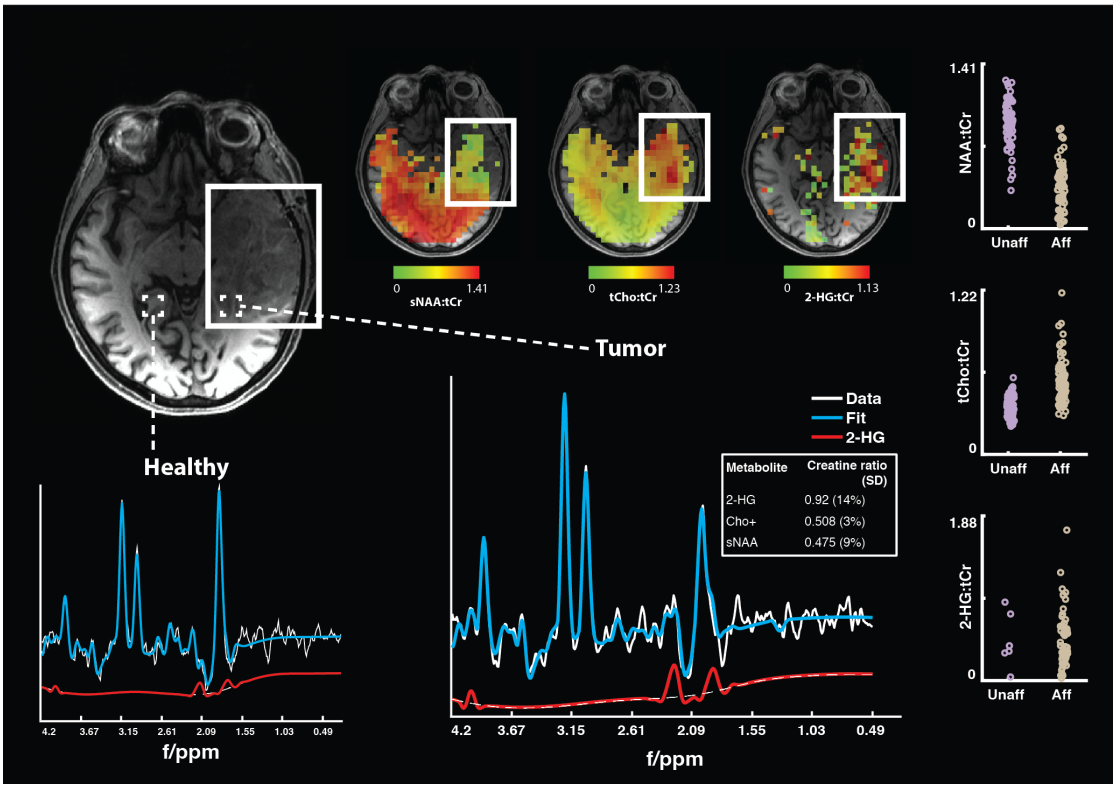


Figure 2.14. High-resolution metabolite mapping of a high-grade oligodendroglioma in a 52-year-old female. Regional increase in NAA and decrease in tCho in the affected area is consistent with tumour pathology. In addition, 2-hydroxyglutarate, an important marker of malignancy, is present within the tumour, but virtually absent in the unaffected hemisphere. Example spectra from unaffected (left) and affected (right) tissue are shown in the lower panels. Figure from [204].

2.3.1 Magnetic resonance spectroscopy: the basics

MRS relies on the MR properties as other MRI modalities. Consistent with the discussion of MR above, here I specifically discuss ^{1}H MRS, but MRS for other molecules is also possible.

^{1}H MRS measurements quantify the molecules bound to protons and quantifies them based on the oscillation frequency of the FID.

The oscillation frequency of the FID depends on two factors, chemical-shift and J-coupling [205, 206]. Chemical-shift refers to the change in the FID due to the electrochemical environment surrounding the atomic nucleus. When bound to a compound, the ^{1}H molecules may be shielded from the B_0 field by the surrounding

electrons that shifts the FID frequency. The specific frequency shift depends on the compound to which the proton is bound.

J-coupling results from the interaction between neighbouring atomic nuclei within a compound depending on their configuration (i.e., the geometry of the molecule, types of ionic bonds, etc). Again, the frequency shift resulting from J-coupling can be predicted if the molecular structure is known (and, in a pure sample, J-coupling can be used to infer the geometry of the compound).

Together, chemical shift and J-coupling provide the information necessary to infer the composition of the tissue being studied.

2.3.2 Magnetic resonance spectroscopy: localization

The localization method is one key feature that defines spectroscopic acquisitions. Two of the most common localization methods are Point RESolved Spectroscopy (PRESS; [207]) and STimulated Echo Acquisition Mode (STEAM; [208]). These acquisitions use three RF pulses with the slice-selective gradients (x-, y-, and z-gradients) active (e.g., PRESS: 90-degree excitation to select along the z-axis, followed by two 180 degree re-focusing pulses to select x- and y-axes); thus, the voxel volume is defined by the intersection of the three slices.

Adiabatic pulses for excitation provide superior signal to diabiatic versions of PRESS and STEAM. Rather than a distinct excitation pulse that flips M_0, adiabatic pulses slowly ramp RF power tip M_0 to the desired angle. This approach achieves more homogenous excitation, yielding higher signal as M_0 relaxes. Adiabatic pulses are also less sensitive to spatial inhomogeneity in the B_0 field and can be applied with higher RF power, thereby achieving greater SNR than diabiatic sequences (e.g., PRESS). However, the greater amount of RF energy absorbed by the participant (specific absorption rate; SAR) in the scanner can cause adverse

effects. The MRSI sequence used a partially-adiabatic implementation of the PRESS localization (semi-Localization by Adiabatic Selective Refocusing [semi-LASER]; [209]).

2.3.3 Magnetic resonance spectroscopic imaging

MRSI, on the other hand, permits multiple MRS voxels to be acquired simultaneously. MRSI uses single-voxel MRS (svMRS) pulses to localize a volume of interest, which is subdivided though phase-encoding (e.g., EPI k-space sampling). MRSI approaches are available for both 2D [204, 210, 211] and 3D acquisitions [212]. In the present thesis, I use a 2D acquisition.

2.3.4 K-space sampling

For MRSI, the most common k-space sampling method is linear echo-planar spectroscopic imaging (EPSI; [213]). However, in the present thesis, I use a non-linear (circular) trajectory. Non-linear approaches trajectories are more efficient than linear sampling methods, and, of these methods, circular trajectories are less prone to gradient imperfections [214]. Thus, there is substantial support for the use of concentric circular trajectories, or concentric ring trajectories (CRTs), for MRSI acquisition.

Three parameters define CRT: the number of rings (N_{rings}), the number of times each ring is acquired ($N_{averages}$), and the diameter of the outer ring (α), which defines the width of the spatial response function (SRF) main lobe (Figure 2.15). To acquire the rings, the readout gradients x- and y- are played out in sine and cosine during the FID.

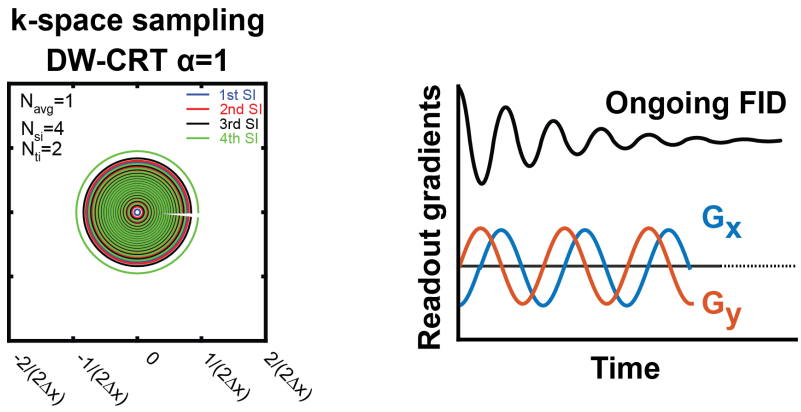


Figure 2.15. K-space sampling with concentric ring trajectory. Left. Example DW-CRT readout. Right. K-space sampling is accomplished by applying the readout gradients Gx and Gy in sine/cosine manner to sample concentric rings (N=24) of varying radii. With density weighting, rather than acquire multiple samples of the same 24 rings, the rings are spaced to optimally shape the spatial response function.

Density-weighted (DW) ring spacing improves CRT. Rather than acquiring several repetitions of each ring with radially equidistant (RE) spacing, using DW-CRT, the rings are spaced such that, for a given number of rings, k-space is maximally sampled based on the Hamming distance:

$$w(k) = \frac{\Delta x}{2}(1 + cos(2\Pi k \Delta x / a))$$

Where Δx represents the nominal spatial resolution. The desired density, *w(k)*, is calculated by integrating w(k) from 0 to the k locations corresponding to equal area bins taken as the radii for the sampling rings [210]. α defines width of the outer ring, and, importantly, as α increases to a value greater than 1, the intermediate regions of k-space are sampled more coarsely. At higher α, the main SRF lobe is well defined but the coarse sampling causes non-linear blurring of signal to distant voxels (Figure 2.16). When α equals 1, the rings are optimally distributed between $-1/(2\Delta x)$ and $1/(2\Delta x)$.

The difference in the SRF (Figure 2.16), effective SNR and linewidth (LW) of the resulting spectra (Figure 2.17) between DW-CRT with alpha = 1 to EPSI, RE-CRT, and DW-CRT alpha = 1.71 are shown below.

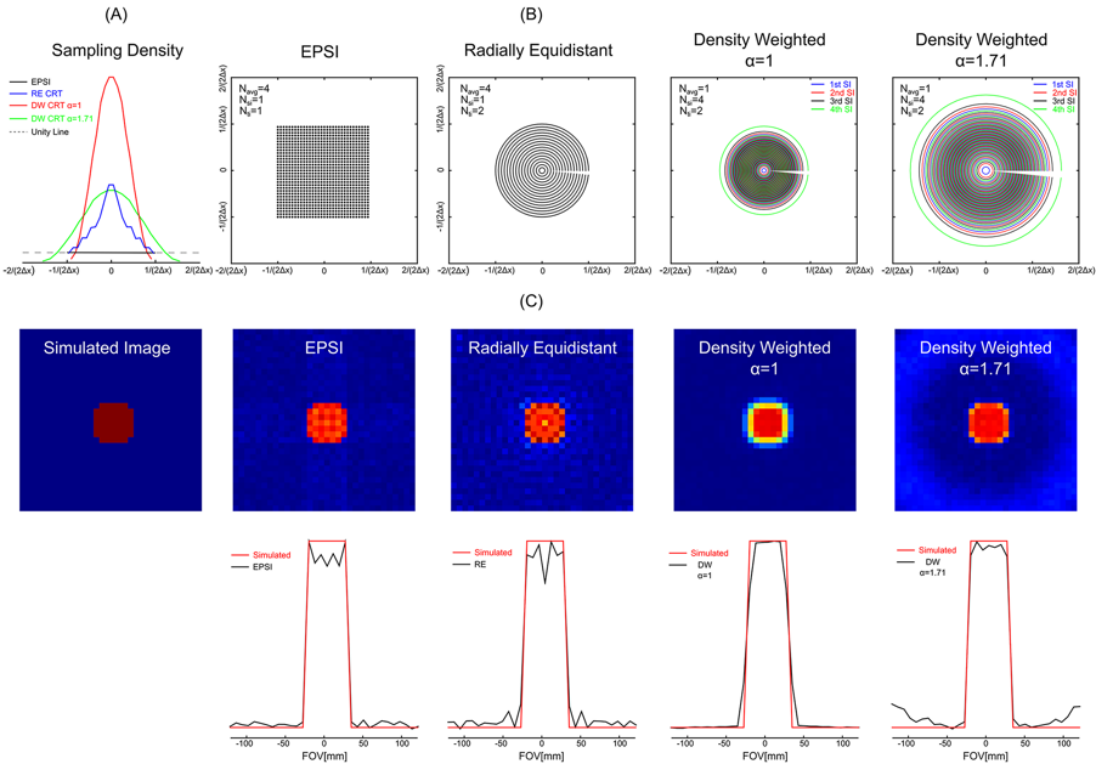


Figure 2.16. Spatial response function (SRF) across k-space sampling methods. Top. K-space sampling patterns of echo planar spectroscopic imaging (EPSI), radially equidistant concentric rings, density weighted concentric rings [equivalent to the present study], and density weighted concentric rings with increased ring spacing. Middle, simulated image and resulting reconstruction using each k-space sampling pattern. Density weighted concentric rings with alpha = 1 provides the most punctate and faithful reconstruction of the image with substantially less blurring between voxels. Lower. One-dimensional projection of the simulated SRF of each k-space sampling strategy. The long tails of the SRF result in image blurring. Density weighted concentric rings with alpha = 1 most closely matches the ideal shape. Figure reproduced from [210].

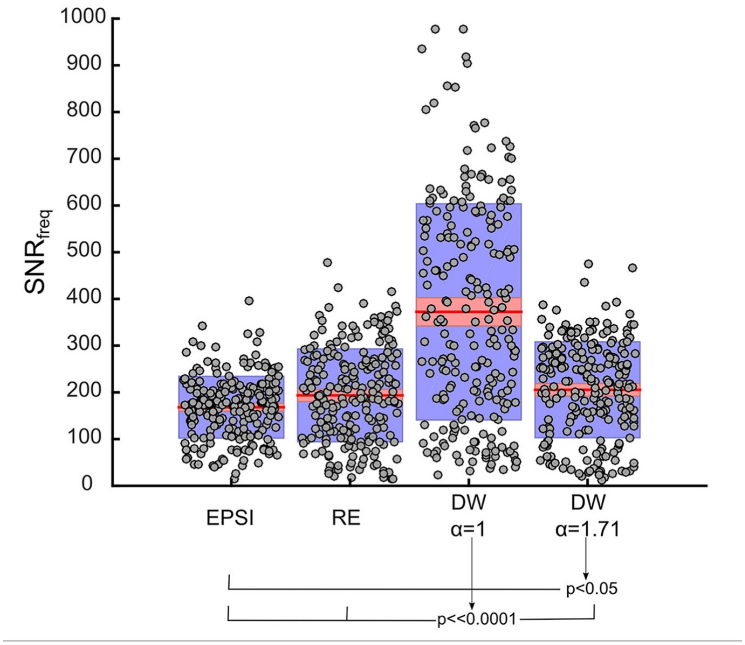


Figure 2.17. Spectral SNR of MRSI measurements from a phantom. MRSI was acquired with echo planar (EPSI), radially equidistant concentric rings (RE), density-weighted concentric rings with alpha = 1 (DW alpha =1, used in the present thesis), or DW with alpha = 1.71. DW with alpha = 1 achieved higher signal to noise than any other technique. Each point represents a voxel in the image. Red line indicates mean, with standard error of the mean (SEM) shown in red and standard deviation (Blue). Figure reproduced from [210].

2.3.5 Metabolite cycling for noise correction

Metabolite-cycling provides critical improvements in spectral quality. The details of metabolite-cycling are described in detail below.

For background, water is the molecule in highest abundance in the human brain: water concentration is roughly in the tens of moles, while other neurochemicals are present in the milli- and micro-molar range [215]. Therefore, the contribution of water must be removed to enable quantification of other neurochemicals with proton MRS. Water-suppression pulses are often incomplete, which leads to distortions in the spectra that prevent quantification. Metabolite-cycling is an acquisition and processing technique whereby spectroscopic measurements are acquired without water suppression.

For metabolite-cycling, during acquisition, excitation pulses flip the metabolite spectra upfield and downfield relative to the water peak in an interleaved fashion. The water concentration is computed by subtracting the upfield and downfield signals, which leaves the metabolite spectra as a residual (Figure 2.18). In the present thesis, I implement metabolite-cycling with semi-LASER MRSI [204].

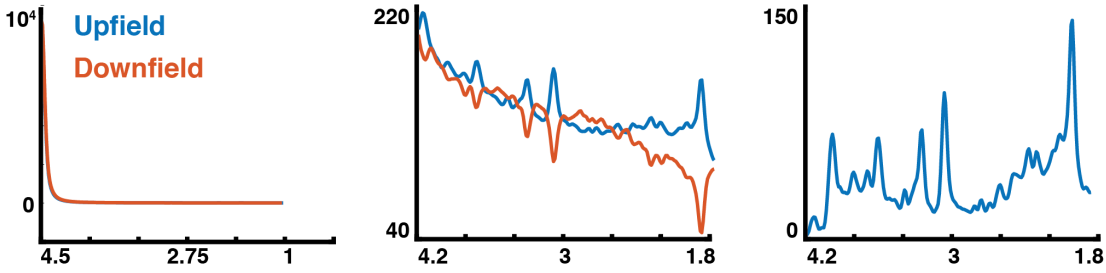


Figure 2.18. Diagram of non-water suppressed MRS acquisition. Left. Fourier transform of the acquired FID, both upfield and downfield spectra show a large peak corresponding to the resonant frequency of water. Middle. Zoomed view of the upfield and downfield non-water suppressed spectra showing additional metabolite peaks. The subtraction of these peaks yields a spectrum with the water signal removed (right).

Metabolite-cycling also allows frequency alignment to correct motion and hardware-related artefacts [216], and eddy-current correction to correct for gradient imperfections during pulse localization [211].

2.3.6 Further preprocessing of MRSI

Applying the Hankel-Lanczos singular value decomposition algorithm (HLSVD; [217, 218]) to fit and remove any remaining water signal greatly improves data quality. HLSVD optimizes the fit of n exponentially dampened sinusoids to the FID in the time domain based on the single-value decomposition of a Hankel matrix formed from the FID. The fitted sinusoids that fall within the water frequency range are subtracted from the data. The success of this method depends on size of the Hankel matrix extracted from the FID and the number of

sinusoids fit to the matrix [218]. Here, the Hankel matrix size was based on the study by Cabanes and colleagues, and number of sinusoids fitted to the data were determined by minimizing the resulting root mean squared error of the model fit [218].

In addition, large macromolecules such as proteins and lipids, found in high concentrations in the meninges, cause distortions in the spectral signal. For this reason, svMRS acquisitions generally avoid the meninges. This solution is suitable for MRSI, as well. Outer volume suppression (OVS) pulses also solve the lipid problem at the acquisition stage, but OVS requires substantial investment of experimental time for frequency calibration.

During preprocessing, a model of the lipid signals the spectroscopic image can be derived, and subsequently removed from, the data itself via L2-normalization [204, 219]. L2-normalization can remove the macromolecule signal, permitting quantification of clinically-relevant metabolites that occur in high concentrations (e.g., NAA, Cho, and Cr). However, neurochemicals in small concentration, such as GABA, cannot be reliably quantified after L2-normalization. Thus, I chose to acquire data while avoiding meninges, rather than use OVS pulses or L2-normalization.

2.3.7 Metabolite quantification

After preprocessing, the metabolites in each voxel are quantified. Quantification is accomplished by fitting simulated spectra from a set of molecules known to occur in brain tissue, i.e., the basis set, to the acquired data (Figure 2.19). In this thesis, I implement this procedure with the LCModel software [220].

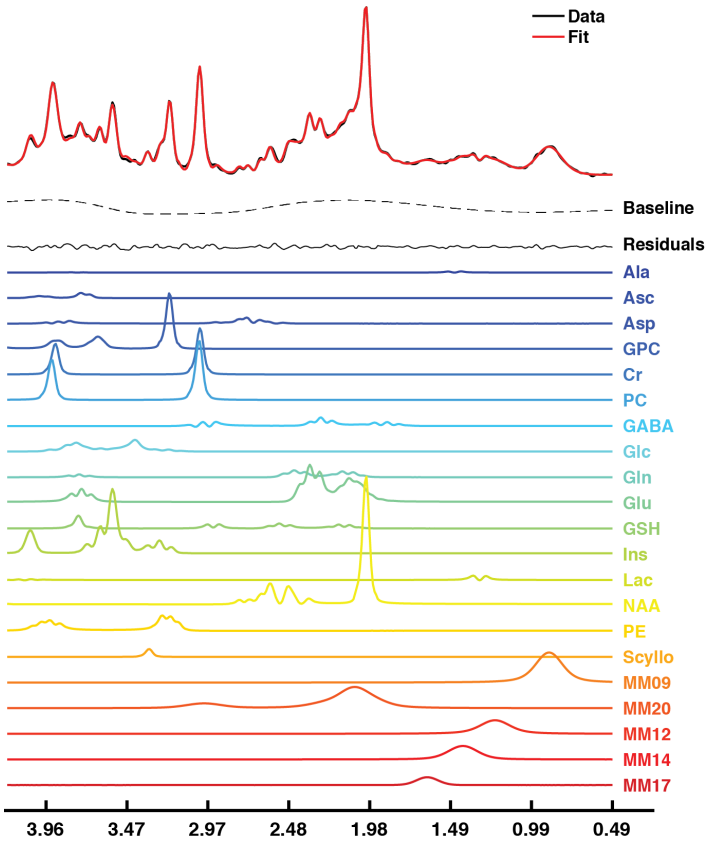


Figure 2.19. Example basis set and LCModel fit from a single MRSI voxel.

While obtaining absolute metabolite concentrations using MRS is possible, it requires assumptions of B_0 homogeneity and consistent SNR across scans, which are often violated. Therefore, it is typical to quantify neurochemical concentrations relative to a reference metabolite chosen based on its stability and robustness. I used total creatine (tCr), the sum of creatine and phosphocreatine, for this purpose.

The estimated lowest unbiased parameter-estimate of the fit (Cramer-Rao Lower Bound; CRLB) can be used assess the quality of the neurochemical measurement [221]. However, relative CRLB values (expressed in percentages) are often higher when metabolite concentration is low, and CLRB values may appear low even if data quality is poor [222]. For this reason, in this thesis, I used visual inspection, CRLB, line-width, and SNR to assess data quality.

The close proximity of the GABA spectrum to other high concentration metabolites such as NAA and creatine makes quantification of GABA at 3-Tesla particularly challenging. Spectral editing (e.g., MEscher-GArwood Point RESolved Spectroscopy, MEGA-PRESS), where additional pulses are used to invert specific neurochemical contributions, can be used at 3-tesla to improve GABA measurements [223]. However, editing approaches rely on subtraction of serially acquired scans, which can be sensitive to system imperfections and participant motion. Therefore, I used the aforementioned semi-LASER acquisition.

Finally, because MRS measures of neurochemicals reflect the concentration across the whole voxel, whether the GABA measured with MRS is located in the synapse, vesicles, glia, or in peri-cellular space is not known. As such, linking the MRS-measured GABA to synaptic activity is not straightforward. MRS measurements of GABA concentration in primary motor cortex are not related to TMS measures of $GABA_A$ or $GABA_B$ receptor activity, but rather are related to a TMS measure of GABAergic tone [224]. The relationship between GABAergic tone and synaptic function is unclear. However, studies of functional properties, such as cortical tuning in somatosensory cortex [225] suggest that GABAergic tone plays a role in establishing local excitability and regional cortical properties.

3 Evidence for task-specific functional plasticity in humans

Resting-state fMRI is well-suited for investigating functional plasticity that occurs as a consequence of training. However, whether this technique is sensitive to task-specific changes, as opposed to non-specific, domain-general changes, has not been systematically investigated. In the present study, I assess the sensitivity of resting-state functional connectivity analyses to task-specific effects. I investigated this question with a dataset that used a longitudinal design that included two different tasks (a visuo-spatial learning task and a motor-sequencing task) and non-training control visits. I found that the visuo-spatial and motor-sequencing tasks produce dissociable changes in functional connectivity within task-relevant brain networks. This work provides a quantitative assessment of training dependent changes in functional connectivity with a rigorous experimental design, as well as a speculative discussion regarding the origin of non-specific changes in resting-state functional connectivity that occur across the course of the day.

3.1 Introduction

As discussed in Chapter 1, an increasing body of evidence suggests that altering the functional interactions between brain regions may be critical for short-term skill-learning [141]. Such changes can be observed in the period immediately after skill learning in humans using resting-state fMRI. Resting-state MRI has been used to investigate offline training-related change in a wide range

of tasks [147, 150, 174, 226, 227]. However, determining whether training *caused* the changes in resting-state functional connectivity can be challenging.

Evidence for a training-related change in MRI measures is generally considered compelling if the change is demonstrably specific to the training (compared to non-training control) and to the specific task being trained (i.e., task A versus task B) [228]. Correlations between the degree functional connectivity change, and behavioural improvement can further strengthen the case for functional relevance. Typically, MRI studies investigating the effect of training use a longitudinal design in which participants are scanned at baseline, provided training in a task of interest, and then scanned again at the end of the training regimen. While a longitudinal design can provide the strongest support for a causal link between training and brain plasticity, changes in functional MRI measures, such as resting-state functional connectivity or structural MRI measures such as cortical thickness [229], can arise for non-specific reasons. These non-specific effects pose problems when studying brain plasticity.

The lack of an experimental manipulation makes confounding non-specific effects particularly insidious when considering spontaneous brain activity and determining activity that was caused by training as opposed to general memory formation or other reasons may not be possible without rigorous controls. For example, resting-state functional connectivity of the default mode network changes in several different training contexts (e.g., [230, 231]) as well as over the course of the day with no training [232]. Disambiguating general from specific effects in this case is challenging without jointly considering both active (i.e., task of interest) and control tasks in the same participants.

To my knowledge, task-specificity of resting-state functional connectivity change has not been directly tested using multiple controls. Thus, in the present study, I undertook a comprehensive exploration of task-specific changes in resting-state functional connectivity. I investigated this question using a dataset collected with a longitudinal design that featured training on both visuo-spatial and motor-sequencing skills. First, I investigated whether training-related changes in resting-state functional connectivity were specific to the task being trained. Second, I examined the relationship between task-specific changes in resting-state functional connectivity and improvements in task performance after training. Finally, I examined non-specific effects that occurred even in the absence of training.

3.2 Methods

3.2.1 Participants

Twenty-one adult volunteers participated in the present study. One participant was removed due to scheduling constraints, leaving 20 participants (mean age: 25.8 years; range: 20–38, 10 females). All participants were screened for any history of neurological and psychiatric disorders, as well as drug and alcohol abuse. In addition, participants were excluded from the study if they played racing videogames and if they had extensive musical training using instruments that require fine movements of the digits. Participants were right-hand dominant based on the Edinburgh handedness inventory (Oldfield, 1971). All participants gave written informed consent. The consent and protocol used for this study were approved by the NIH Institutional Review Board (Protocol 93-M-0170, NCT00001360).

3.2.2 Study design

As illustrated in Figure 3.1, each participant was scanned once each week, over four consecutive weeks (Visits). During each Visit, participants completed two MRI sessions, i) 1000 to 1200 (AM session) and ii) 1400 to 1600 (PM session). There was a two-hour gap between scan sessions during which participants experienced training or non-training. During Visits 1 and 4 (non-training), participants were given no specific tasks during the two-hour gap. For Visits 2 and 3 (training), in the time period between the AM and PM sessions, participants completed computer-based tasks (Visit 2: visuo-spatial, Visit 3: motor-sequence learning). In addition, to investigate the effect of long-term training on structural MRI measures, between Visits 3 and 4 participants completed additional training on the motor-sequence task. However, here I focus only on the impact of short-term training on resting-state functional connectivity. For pragmatic reasons the order of the visuo-spatial and motor-sequence tasks were not counterbalanced. During all imaging visits, the timing of the scans was strictly observed, with minor deviations due to administrative and technical issues during a small number of scanning sessions.

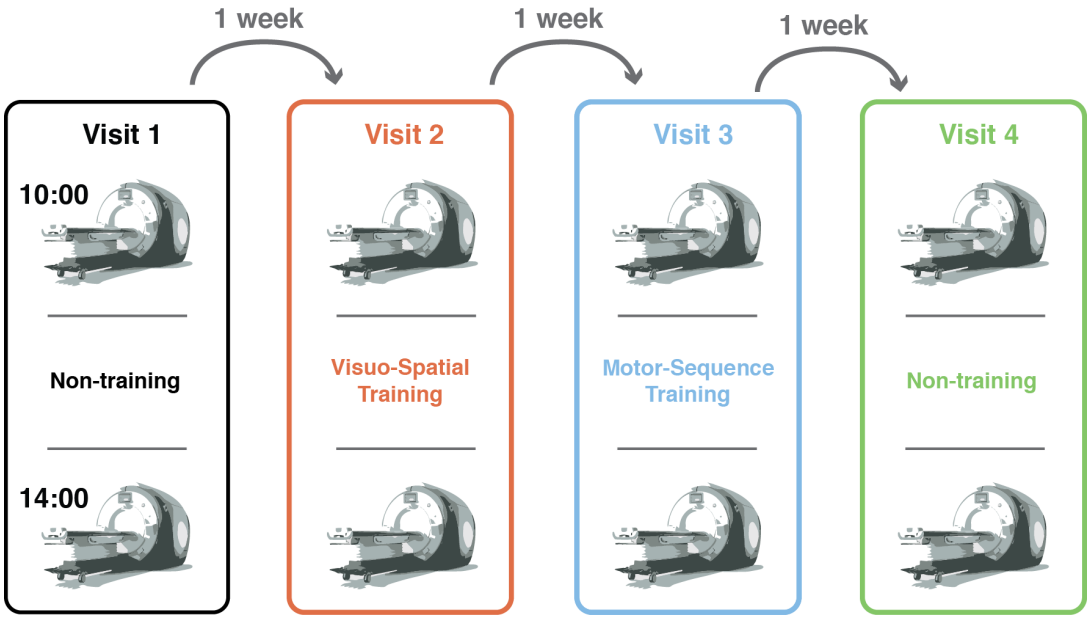


Figure 3.1. Experimental design. Participants were scanned over 4 separate visits, a week apart. During each Visit, participants were scanned from (approximately) 10:00-12:00 (AM) and 14:00-16:00 (PM). On Visits 1 and 4 participants were not given any training bet ween the AM and PM scans, participants were provided training on either a visuo-spatial task or a motor-sequence learning task, respectively.

3.2.3 Behavioral tasks

To investigate task-specific effects of training, I used behavioural paradigms (visuo-spatial and motor-sequence learning) that are thought to involve dissociable neural systems. The visuo-spatial task involved learning the spatial layout of a racetrack in a computer-based videogame, and the motor-sequencing task involved pressing a series of buttons as quickly as possible with the non-dominant hand. The two tasks are described in detail below.

3.2.3.1 Visuo-spatial task

During Visit 2, immediately following the AM scan, participants underwent 90 minutes of training to learn the spatial layout of a racetrack while playing a racing videogame. Training was provided on a 20-inch LCD monitor and participants navigated the racetrack with their right (dominant) hand using the

arrow keys on a standard computer keyboard. The racing videogame (Need for Speed, Porsche 2000 Edition; Electronic Arts) was chosen to match a previous study wherein it was reported that two hours of training evoked structural changes in the hippocampus and parahippocampal cortices, regions that are known to be crucial for spatial navigation [233]. Therefore, the training protocol used in this experiment followed the protocol described in the prior study as closely as possible.

Briefly, participants were instructed to learn the spatial layout of the racetrack (Monte Carlo circuit) and improve their lap times. In total, participants completed 20 laps, grouped into 5 blocks of 4 laps each. Before the onset of training (i.e., baseline), and after each of the 5 training blocks, participants were shown a series of 12 randomly arranged screenshots of the racetrack and were instructed to rearrange the screenshots in the order they deemed sequential. In addition, participants were instructed to draw a bird's eye view of the racetrack on a sheet of paper. The initial lap time and image arrangement were recorded, and these data were considered as the baseline for each participant. Image arrangement was considered the primary outcome measure for the visuo-spatial task, as improvement in this measure is more specifically related to spatial memory (lap time, which reflects both motor and spatial memory improvements).

3.2.3.2 Motor sequence learning

During Visit 3, participants completed an explicit sequence learning task. During the inter-scan period, participants underwent Baseline, Training, and Test epochs. The baseline epoch was used as the basis for estimating motor-skill improvement after training. On each trial, participants were shown a series of

eight items (digits 1-4) on a screen. The participants were instructed to use the index, middle, ring, and little finger to enter the eight-item sequence as fast and accurately as possible on a cylindrical button box (Current Designs, Cambridge, MA). After completing each eight-item sequence, the participant pressed the thumb button to indicate completion; auditory feedback was provided after the thumb-press based on accuracy (bell = correct, beep = incorrect). In the baseline epoch, participants were assigned two different 8-item sequences, which were displayed in a counterbalanced fashion over 6 blocks of trials (12 trials per block). Both left and right hand were tested over the 6 blocks in an alternating fashion. Initial motor skill was assessed by comparing performance on each sequence with each hand (18 trials per hand/sequence pairing).

After the baseline epoch, one of the sequences was randomly chosen to be the "trained" sequence. During the training epoch, participants were instructed to respond using only the left (non-dominant) hand and to be mindful of the number to digit mapping. On each trial, participants were shown the "trained" eight-item sequence and given 1 minute to enter the sequence as many times as possible using the buttons mapped to the respective fingers and indicate completion by the pressing the thumb button. Each 1-minute period was considered a single trial, and trials were grouped into blocks of 5 trials. Performance was indexed by counting the total number of correctly entered sequences and the average time required to correctly input the sequence. The training session was self-paced, and participants were encouraged to take a short break after each trial to reduce muscle fatigue. After completion of a trial, the next trial was initiated by pressing the space bar.

The trained sequence was continuously presented over 7 blocks of 5 trials per block (total blocks = 35). Participants continued to receive auditory feedback throughout the trials. To encourage continued improvement, after each trial participants were shown the number of correct sequences completed and the average RT required to enter the correct sequence.

Immediately following the training epoch, participants started the test epoch, in which short-term sequence learning was assessed. The test epoch followed the same presentation format as the baseline epoch and used the same sequences (1 trained/1 untrained) shown during the baseline epoch. Motor-sequence learning was assessed by comparing performance on each sequence (Trained/Untrained) with each hand (18 trials per hand/sequence pairing). In total, the short-term training assessment (Baseline/Train/Test epochs) lasted approximately 90 minutes. All stimulus presentation was conducted in a room adjacent to the scanning facility using a MacBook Pro with PsychoPy presentation software [234].

Although the focus of the current Chapter is on short-term training, additional training was conducted between Visits 3 and 4 to assess the impact of long-term training on MRI measures. In addition, during Visit 4, before the AM scan session, participants performed the long-term training test following the same procedure described for baseline assessment. This was the only visit where a behavioural assessment preceded the AM scan.

3.2.4 Behavioural data analysis

The assessment of the visuo-spatial task focused on the improvement of memory for the spatial layout of the race track as a function of training, while the motor-sequence task focused on the reduction in RT to execute the trained

sequence. The summary statistics used (e.g., slope of learning or median RT) reliably capture participants' performance improvement, and the behavioural data were analysed using statistical methods that are appropriate given the nature of the behavioural improvement being assessed.

3.2.4.1 Visuo-spatial task

The scene rearrangement test was analysed using methods similar to those described in [233].The slope of the reduction in accuracy for scene arrangement over the blocks of trials was computed using a linear mixed effects model (LME) model implemented in the lme4 package in R [235]. For the behavioural measure of interest (i.e., accuracy) the model consisted of fixed effects for both intercept and slope and a random effect to account for individual variations. The participant-specific slope generated by the LME analysis was used to examine the relationship between behavioural performance and training-specific changes in resting-state functional connectivity.

3.2.4.2 Motor-sequence learning

To demonstrate that sequence learning is specific to the sequence type and trained hand, a repeated measures analysis of variance (ANOVA) was implemented in SPSS, modelling the main effects and interaction between three factors: Timepoint (Pre/Post training), Hand (Right/Left), and Sequence Type (Trained/Untrained). Thus, a significant three-way interaction between Timepoint, Hand, and Sequence type would serve as evidence that performance improvement is specific to the sequence type and trained hand.

3.2.5 MRI acquisition

Participants were scanned on a 3 Tesla General Electric MR750 scanner with a GE 32-channel head coil. Multimodal MRI data were acquired during each scan session. Of relevance to the present analysis are the functional and anatomical MRI scans, and these are described below.

3.2.5.1 FMRI resting-state scan protocol

Two 5-minute resting-state scans were acquired during each session, giving a total of 10-minutes of resting-state data per scan session (both AM and PM) for each participant. [EPI scan parameters: sequence: 2D-EPI, 80 × 80 matrix, slice orientation: axial; Coverage: whole-brain, Number of slices: 43; resolution: 3 x 3 mm; Slice thickness: 3 mm; TR: 2000 ms; TE: 27 ms; flip angle: 77˚; acquisition time: 5 min]. Respiration data was acquired using the scanner's acquisition system. During each session, participants were instructed to rest with eyes open while fixating on a cross. A qualitative assessment for head motion was performed using real-time motion parameters extracted during the scan and if there was any evidence for significant head motion (i.e., > 2mm), the scan was repeated immediately after the failed scan.

3.2.5.2 FMRI motor-sequence localizer

To localize sensorimotor cortex, a functional localizer scan was acquired during Visits 1 and 4 from all participants. The localizer scan was always performed at the end of the scanning session, and never occurred before the resting-state scan. Participants were counter balanced so that half the participants performed the task during the AM session and the other half performed it during the PM session. Scan Parameters: 2D-EPI sequence (80 × 80 matrix size, slice

orientation: axial; Coverage: whole-brain, Number of slices: 43; resolution: 3 x 3 mm; Slice thickness: 3 mm; TR: 2000 ms; TE: 27 ms; flip angle: 77°; acquisition time: 6 min 40 s).

The MRI-safe button box that was used for training was placed on the participant's abdomen. Prior to scanning, participants were given sufficient practice to perform the task using both the right and left hands while in the scanner and to familiarize themselves with the button box. The trials were presented in an on/off design with alternating blocks of fixation trials and motor-sequence trials. During each 16 second block, participants viewed a fixation screen or one of two motor-sequences that was randomly selected for each participant. During the fixation trials, eight zeros were presented to the participant, and they were instructed to keep their hands idle and focus on the screen. During the motor-sequence trials, each 8-item sequence was presented with a letter indicating the hand to be used for entering the sequence. Participants were instructed to enter the sequence in a self-paced manner until they saw the fixation screen. The localizer experiment comprised 12 sequence blocks. Both the sequence (trained or untrained) and the order of left and right hand blocks was randomized for each subject. The entire localizer scan lasted 6 minutes and 40 seconds.

3.2.5.3 T1-weighted anatomical scan

Two separate T_1weighted MPRAGE volumes were acquired during each session. [scan parameters: sequence: accelerated sagittal 3D MPRAGE, FOV: 25 cm, matrix size: 256 × 256, slice thickness: 1 mm; TR: 7 ms; TE: 3.42 ms; flip angle: 7°; SENSE acceleration factor (ASSET): 2; acquisition time: 5 min 23 s].

3.2.6 FMRI data pre-processing

All EPI data (localizer and resting-state) were preprocessed using AFNI [198].

3.2.6.1 Motor sequence localizer preprocessing

The first and last 8 TRs from each scan were removed. Signal outliers were attenuated [192] (3dDespike) and slice time correction was applied (3dTshift). Scans were aligned to the first image in the time series (3dVolReg) and registered to the MNI-152 template. The data were smoothed with a 3 mm Full-width half-maximum (FWHM) Gaussian kernel.

3.2.6.2 Resting state preprocessing

After image reconstruction, respiration artefacts were removed on a slice-wise basis with RETROICOR correction procedure [191]. The first 3 images from each scan were removed. Signal outliers were attenuated [192] and slice-time correction was applied.

Motion correction parameters were estimated relative to the first image. Registration parameters for aligning the EPI to the structural scan for that session were calculated. For alignment across sessions, the structural scan from each Visit was aligned to the mean structural scan estimated from the Freesurfer longitudinal analysis pipeline [236]. These registration parameters (session structural and longitudinal structural) were then combined with the volume-registration parameters and applied in one step (3dAllineate).

Following alignment, I applied the basic ANATICOR [192] procedure to yield nuisance time series for the ventricles and local estimates of the BOLD-signal from the white matter.

All nuisance time series [six parameters of motion, local white matter, ventricle signal, and 3 respiratory noise regressors (AFNI program: RetroTS)] were detrended with fourth order polynomials. These time series, along with a series of sine and cosine functions to remove all frequencies outside the range 0.01-0.25 Hz [237] were regressed out in a single regression step (AFNI program: 3dTproject). Time points with motion greater than 0.3 mm were removed from the data and replaced with values obtained via linear interpolation in time.

The transformation of the mean structural scan from the Freesurfer longitudinal analysis pipeline to the MNI-152 template was calculated (@auto_tlrc), and this transformation was applied to the cleaned data. These data were then smoothed with a 6mm FWHM Gaussian kernel. The two 5-minute resting-state periods were concatenated in time to form a single time series (duration: 10 minutes) from each imaging session for use in the functional connectivity analysis.

The global correlation (GCOR; [238]) and first-difference of motion (@1dDiffMag) were calculated for each scan to include as covariates of no interest in group analysis. Here, GOR refers to the average correlation of all voxels to the rest of the brain and is an effective means of removing variability due to motion and noise factors at the group level. GCOR is distinct from global signal regression, which was not used here [239]. Scan sessions where the first-difference of motion was greater than 0.1 were excluded from group analysis. The number of scans excluded are shown in Table 3.1.

Table 3.1. Inclusion of individual scans based on motion estimates. "X" indicates an included scan, '-' indicates a scan not included.

Subject	S01	S02	S03	S04	S05	S06	S07	S08	Total
P01	X	-	-	-	X	-	X	X	4
P02	X	X	X	X	-	X	X	X	7
P03	X	X	X	X	X	X	X	X	8
P04	X	-	X	-	X	X	X	X	4
P05	X	X	X	X	-	X	X	X	7
P06	X	X	X	X	X	X	X	X	8
P07	X	X	X	X	X	X	X	X	8
P08	X	X	X	X	X	X	X	X	8
P09	X	-	-	X	X	X	X	X	6
P10	X	X	X	X	X	X	X	X	8
P11	X	X	-	X	X	X	X	X	7
P12	X	X	X	X	X	X	-	X	7
P13	X	X	X	X	X	X	X	X	8
P14	X	X	X	X	X	X	X	X	8
P15	X	-	-	X	X	X	X	-	5
P16	X	X	X	X	X	X	X	X	8
P17	X	X	-	X	X	X	X	X	7
P18	X	X	X	X	X	X	X	X	8
P19	X	X	-	X	X	X	-	-	5
P20	X	X	X	X	X	-	-	X	6
Total	19	16	13	18	18	18	17	18	

3.2.7 ROI localization

I focused my analyses primarily on two ROIs that are thought to be selectively involved in learning on these two tasks: bilateral hippocampus and bilateral sensorimotor cortex. Bilateral hippocampus was defined based on the Eickhoff-Zilles maximum probability map atlas included in the AFNI software package (Eickhoff et al., 2006).

Bilateral sensorimotor cortex was defined based on the localizer scan conducted in the first session. Activation maps during the motor localizer scan were calculated by modelling task performance against baseline (3dDeconvolve). Significant functional activation across the group was computed by performing a t-test (3dTtest++) comparing activation left versus right hand activation. The resulting statistical parametric maps were corrected for multiple comparisons using false discovery rate (FDR) ($q < 0.05$). ROIs were generated from these maps for the contrast (Left hand – Right hand) by taking the contiguous clusters of voxels that exceeded the threshold over the left and right sensorimotor cortices.

3.2.8 Resting-state functional connectivity analyses

For all ROIs, the mean time series was extracted, and the whole-brain correlation maps (Pearson's r) from each resting-state MRI scan were then calculated (3dTcorrMap) yielding resting-state functional connectivity maps for bilateral hippocampus and sensorimotor cortex.

3.2.8.1 Effect of training on functional connectivity

For all ROIs, the mean time series was extracted, and the whole-brain correlation maps (Pearson's r) from each resting-state MRI scan were then calculated (3dTcorrMap) yielding resting-state functional connectivity maps for bilateral hippocampus and sensorimotor cortex.

To test for differences in connectivity patterns across visits, the seed-based connectivity maps were submitted to a LME (3dLME; [240]), which is known to be particularly useful for compensating for missing data. The model comprised 3 factors: hemisphere (left/right), Visit (Visits 1-4), and Time-of-Day (AM/PM, which is equivalent to Pre/Post training in the context of the training visits (i.e.,

Visits 2 and 3). These were included as fixed-effects factors for both ROI sets (hippocampus and sensorimotor cortex) separately. First difference of motion and GCOR were included as nuisance covariates.

3.2.8.2 Correlation between functional connectivity change and behaviour

To test for relations between changes in resting-state functional connectivity and improvement in behavioural performance after training, I identified regions that showed a significant Time-of-Day x Visit interaction based on the group-level LME analysis. I calculated the average functional connectivity changes of each region for each participant and performed a Spearman's correlation analysis between the connectivity change and behavioural performance measure. For the analysis involving hippocampus ROI, I used the slope of the improvement in spatial memory as the behavioural correlate. For the analysis involving sensorimotor cortex ROI, I used the improvement in RT for executing the motor-sequence as the behavioural correlate.

In addition, I also tested for a correlation between resting-state functional connectivity changes and behavioural improvement at the whole-brain level on the two different tasks. For each performance metric, connectivity maps from the relevant ROI (visuo-spatial - hippocampus; motor-sequence - sensorimotor cortex) during the relevant Visit (visuo-spatial - Visit 2; motor-sequence - Visit 3) were compared using 3dLME with Time-of-Day (AM/PM equivalent to Pre/Post training) as a factor and performance improvement as a covariate of interest.

3.2.8.3 Leave-one-subject-out-analysis

To characterize the magnitude of the non-specific effects while avoiding the circularity confound [241], I implemented an iterative leave-one-subject-out (LOSO) procedure comprising three steps.

First, I performed the full LME model (Factors: Time-of-Day, Visit, and ROI) with one subject left out to establish the spatial map of brain regions exhibiting the main effect of Time-of-Day and Time-of-Day x Visit interaction. The resulting statistical map was thresholded ($p < 0.01$, $k = 70$), and the procedure was repeated for each subject.

Next, I computed an overlap map across iterations, which was thresholded based on number of iterations showing an effect (25% iterations, $n > 4$) and cluster extent ($k = 70$).

Finally, I extracted the data from the left-out subject within the ROIs as defined by the cluster mask and carried out a full linear mixed effects model using the 'nlme' package in R on the aggregated data [242]. Where appropriate, I conducted follow-up tests for the interaction of Time-of-Day and Visit using a t-test implemented in the 'lsmeans' function in R [243]. This procedure allowed me to compute a relatively independent estimate of the magnitude of the effect of non-specific effects as well as the effects specific to training. I performed this analysis separately for both sensorimotor cortex and hippocampus.

3.3 Results

In this section, I summarize the key behavioural results that provide evidence for training-related changes in performance before elaborating on the

findings pertaining to changes in resting-state functional connectivity between brain regions.

3.3.1 Impact of training on behavioural performance

Across both training tasks, participants showed significant improvement in the respective behavioural measures. In the visuo-spatial training task, over the course of 20 laps, participants showed a clear improvement in their ability to correctly order the pictures of the scenes encountered while navigating the racetrack (Figure 3.2a). Of note, is the striking individual variability in the rate (i.e., the slope) at which they learned the correct order of the scenes (Figure 3.2b).

Similarly, in the motor-sequence task, compared to baseline performance, participants demonstrated marked individual variability in improvement in RT required to execute the 8-digit motor-sequence after short-training (Figure 3.2c).

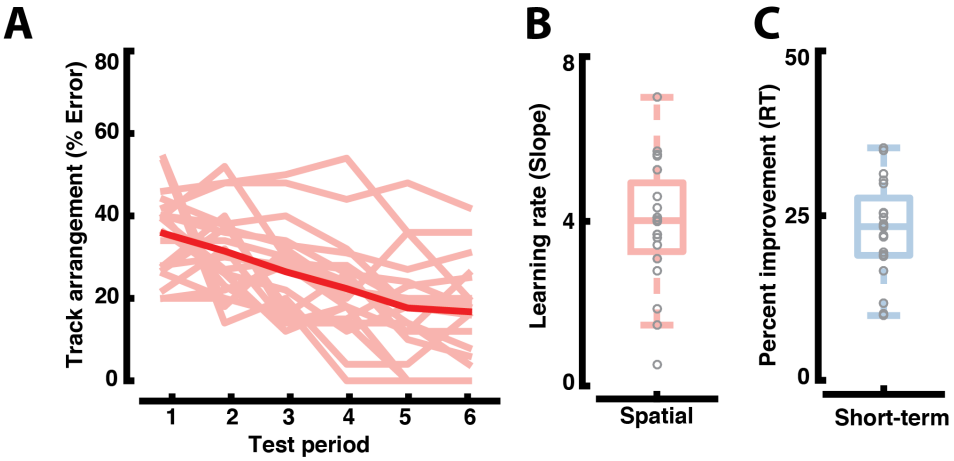


Figure 3.2. Improvement in behavioral performance after training. (A) Participants showed a significant improvement ability to order the scenes of the racetrack over the training period. (B) Visuo-spatial learning rate was highly variable. (C) In the motor-sequence learning task participants showed a significant improvement, in the RT with the trained hand (left hand) after short-term training.

3.3.2 Task-specific effects of training

3.3.2.1 Visuo-spatial training modulates hippocampal functional connectivity

To assess whether these changes were specific to visuo-spatial training, I extracted the estimated difference in resting-state functional connectivity from AM to PM for each region during each visit provided by the LME model (Figure 3.3). Resting-state functional connectivity change was strikingly specific to Visit 2. Compared to the other visits, resting-state functional connectivity decreased from AM to PM between hippocampus and left superior parietal lobe, right dorsolateral prefrontal cortex, left ventral PMC, and right thalamus, during Visit 2. In contrast, resting-state functional connectivity between hippocampus and the posterior cingulate cortex and left putamen increased on Visit 2 compared to other visits. In summary, these results suggest that rapid learning of the spatial layout of the racetrack resulted in task-specific modulation of resting-state functional connectivity with the hippocampus.

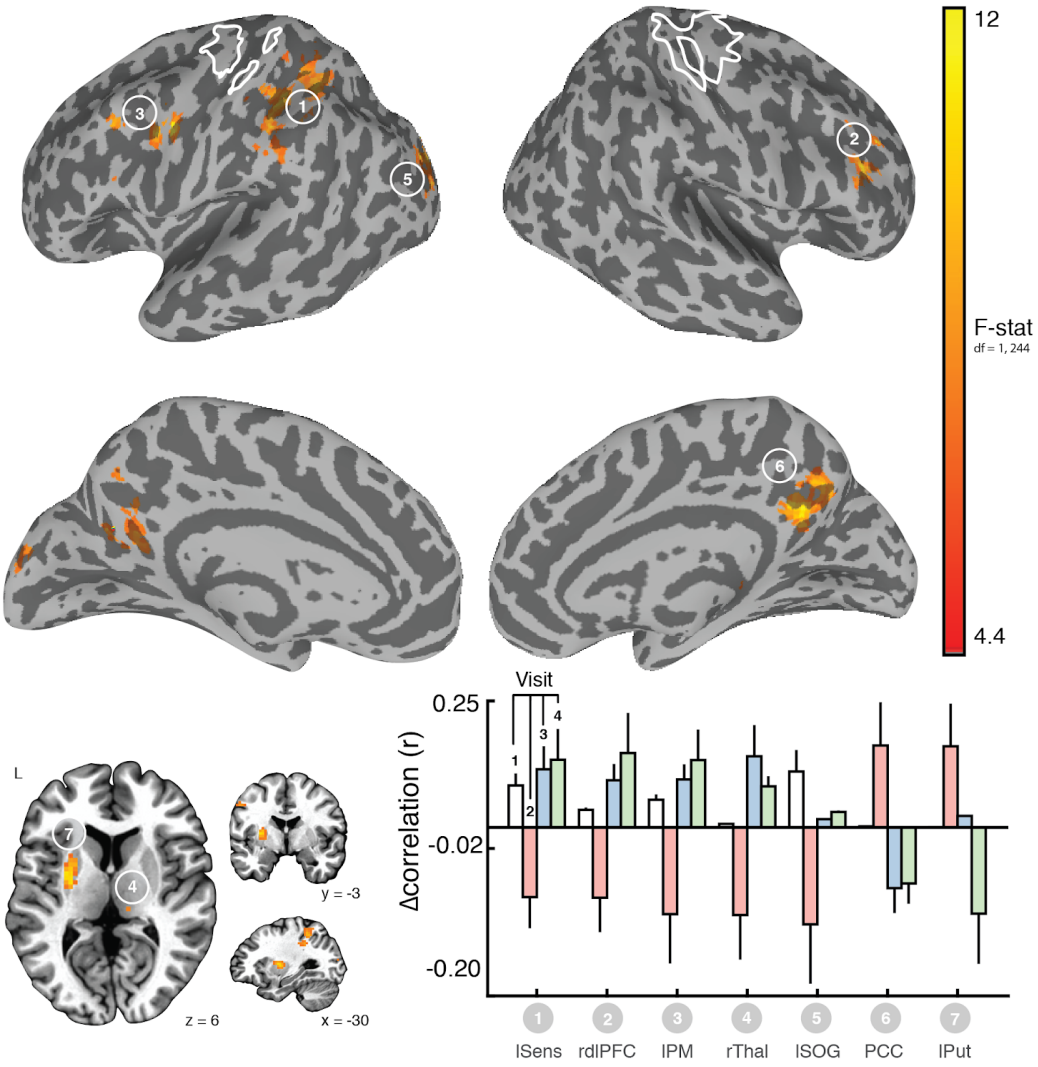


Figure 3.3. Training-specific changes in resting-state functional connectivity with the hippocampus. Statistical parametric maps show the F-statistic for the interaction between Time-of-Day and Visit. The bar graphs show the nature of the interaction in the seven clusters that survived correction for multiple comparisons. The interactions are largely driven by differences in Visit 2 (visuo-spatial training). lSPL: Left superior parietal lobe; rdlPFC: Right dorsolateral prefrontal cortex; lPM: Left PMC; rThal: Right Thalamus; lSOG: Left superior Occipital gyrus; PCC: Posterior Cingulate Cortex; lPut: Left Putamen. Bar graphs are not based on an independent data set and are provided only for the sake of a qualitative assessment of the results.

3.3.2.2 Motor cortex functional connectivity changes after training

Using the same analysis strategy described above, I probed the specificity of training-related changes in resting-state functional connectivity of bilateral sensorimotor cortices. This analysis revealed a significant interaction between Visit and Time-of-Day in bilateral superior temporal gyrus, left lateral occipitotemporal cortex, and right caudate (Figure 3.4).

A qualitative examination of the interactions revealed that they were driven by an increase in resting-state functional connectivity from AM to PM in the bilateral superior temporal gyrus (more prominent on the left superior temporal gyrus) and the left lateral occipitotemporal cortex during Visit 3. Functional connectivity between sensorimotor cortex and the right superior temporal gyrus decreased during Visit 4. In contrast, the interaction in right caudate nucleus appeared to be driven by a strong decrease in resting-state functional connectivity with bilateral sensorimotor cortex during Visit 3, but only modest change in the other three visits. Together, these results provide evidence for modulations in resting-state functional connectivity with the sensorimotor cortices that are spatially distinct from the regions showing resting-state functional connectivity modulations with the hippocampus network but are somewhat more ambiguous with regards to specificity to training.

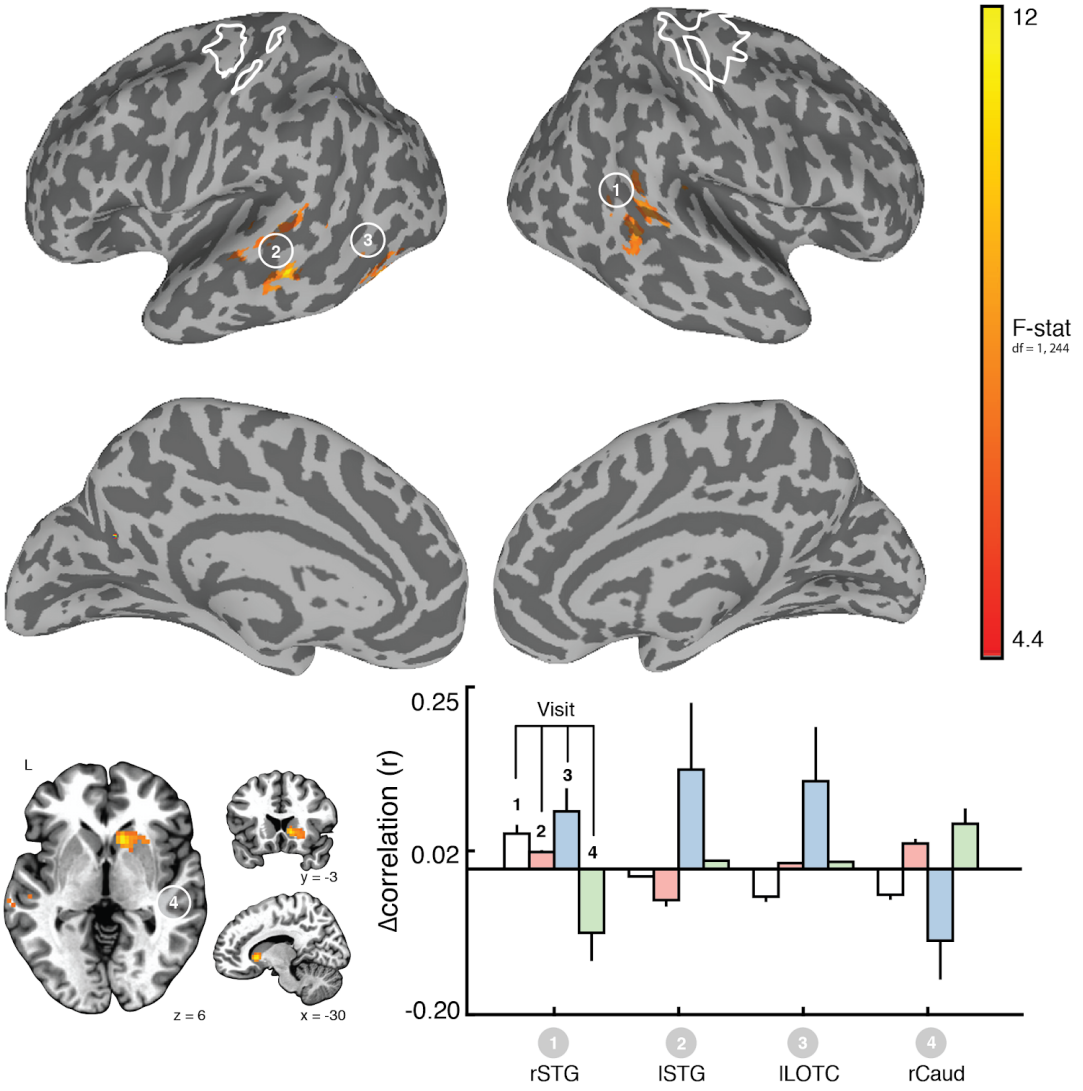


Figure 3.4. Training-specific changes in resting-state functional connectivity with the sensorimotor network. Statistical parametric maps show the F-statistic for the interaction between Time-of-Day and Visit. The bar graphs show the nature of the interaction in the four clusters that survived correction for multiple comparisons. In STG, the interaction appears to reflect differences in Visit 4, while in other regions it appears to reflect differences in Visit 3 (motor-sequence training). rSTG: right Superior Temporal Gyrus; lSTG: Left Superior Temporal Gyrus; lLOTC: Left lateral occipito temporal cortex; rCaud: Right Caudate Nucleus.Bar graphs are not based on an independent data set and are provided only for as a qualitative assessment of the results.

3.3.3 LOSO analysis to determine effect magnitude

To more formally quantify the magnitude of the effects in a non-circular manner, I performed an iterative LOSO procedure to define spatial regions exhibiting the effects of interest based on overlap across the different iterations (see Section 3.2.8.3). Data from the left-out participant was then used to calculate the magnitude of the effect.

The spatial topography of the resting-state functional connectivity changes due to the Time-of-Day X Visit interactions based on the LOSO analysis were very qualitatively similar to the maps observed from the analysis based on all the participants. Overall, the effects observed in the hippocampus were robust: for the Time-of-Day X Visit interaction analysis (Figure 3.5) there was a significant training effect across five of the six ROIs. In contrast, only the caudate showed a Time-of-Day X Visit interaction using the LOSO analysis that was driven by the training visit (Visit 3; Figure 3.6). For both the hippocampus and sensorimotor cortex, the magnitude of the resting-state functional connectivity changes on the visits of interest compared to the other visits was around $r = 0.12$. The non-specific effects, e.g., between sensorimotor cortex and middle temporal gyrus and occipital cortex, was approximately $r = 0.03$.

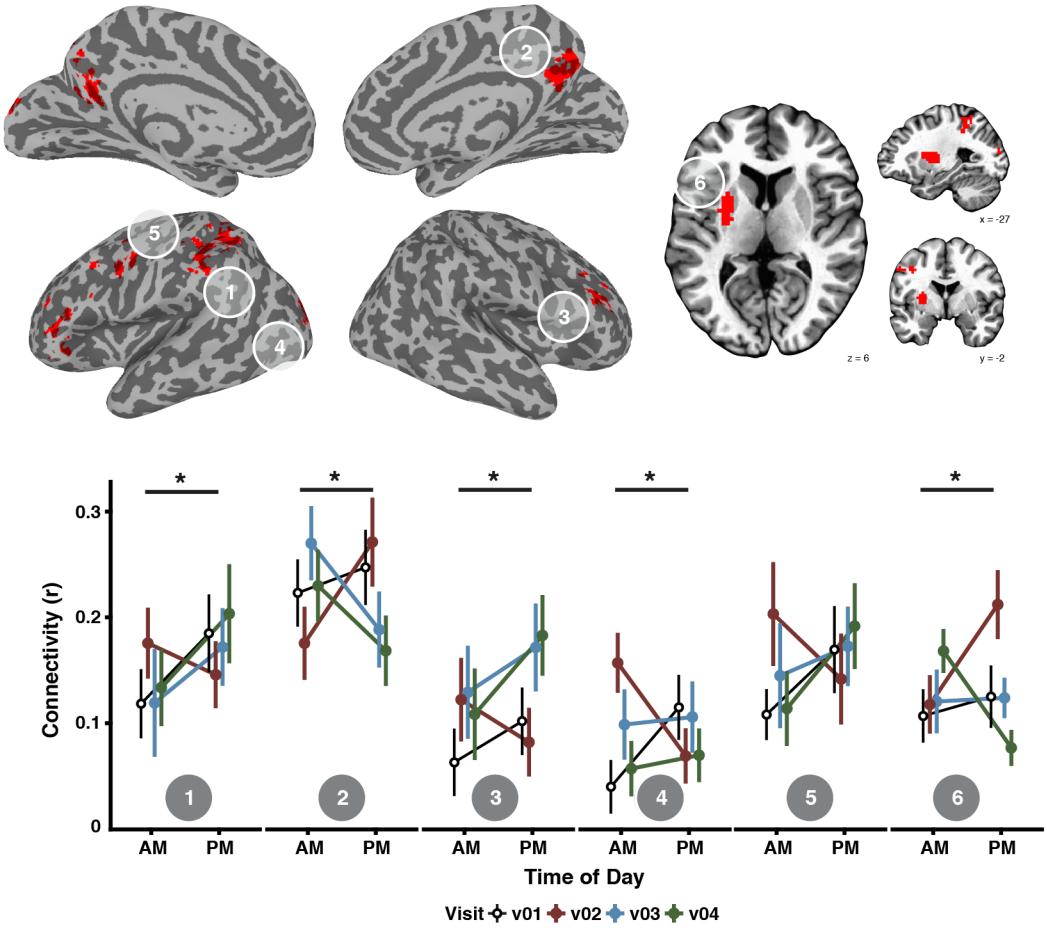


*Figure 3.5. The magnitude of the interaction between Time-of-Day and Visit on resting-state functional connectivity of the Hippocampus. (Upper) The overlap map based on the iterative leave-one-subject-out analysis is generally consistent with the maps generated based on the group analysis Figure 3.3. (Lower) The lineplots show the nature of the interaction and the magnitude of the change in resting-state functional connectivity in each of the ROIs across the 4 visits. * p < 0.05 corrected*

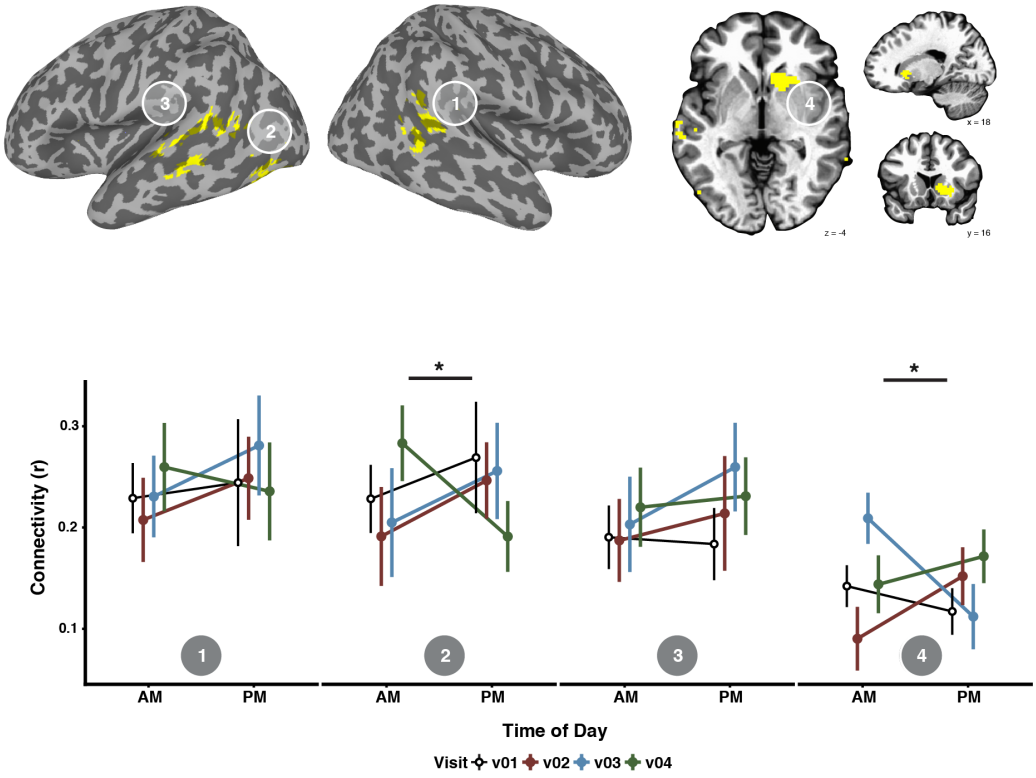


*Figure 3.6. The magnitude of the interaction between Time-of-Day and Visit on resting-state functional connectivity of the Sensorimotor cortex. (Top) The overlap map based on the iterative leave-one-subject-out analysis shows a spatial topography of regions that show a significant interaction effect that is generally consistent with the maps generated based on the group analysis. (Below) The lineplots show the nature of the interaction and the magnitude of the change in resting-state functional connectivity in each of the ROIs across the 4 visits. * p < 0.05 corrected*

3.3.4 Relationship between regions showing functional connectivity change and behaviour

3.3.4.1 Hippocampus ROI

Although the changes I observed were specific to the visit of interest, I wanted to assess whether these task-specific changes are related to behavioural measures of performance improvement. For the visuo-spatial task, I first tested whether changes in resting-state functional connectivity in regions that showed resting-state functional connectivity changes that were specific to Visit 2, correlated with improvements in memory for the spatial layout of the race track. I observed that a change in resting-state functional connectivity between hippocampus and occipital cortex was positively correlated with improvements in visuo-spatial memory performance; however, this correlation did not survive correction for multiple comparisons (corrected $p = 0.05/7 = 0.0071$; $r_{(13)} = 0.68$, $p_{uncorr} = 0.0099$). None of the other regions showed a correlation between change in resting-state functional connectivity with hippocampus and performance in the visuo-spatial memory task.

3.3.4.2 Sensorimotor cortex

For the motor-sequence task, none of the regions exhibiting a Visit x Time-of-Day interaction in resting-state functional connectivity showed a significant correlation with motor-sequence learning.

To further explore the relationship between the resting-state functional connectivity changes and the behavioural changes, I investigated whether any resting-state functional connectivity changes of the hippocampus and

sensorimotor cortex were related to performance outcomes after training on the visuo-spatial and motor-sequence training tasks at the whole-brain level.

3.3.5 Functional connectivity changes reflecting behavioural performance

3.3.5.1 Visuo-spatial learning

Better memory for the spatial layout of the race track correlated positively with increased resting-state functional connectivity from AM to PM between hippocampus and dorsal occipital cortex (Figure 3.7). In addition, this analysis also revealed a negative correlation between improvement in visuo-spatial memory and increases from AM to PM in resting-state functional connectivity between bilateral hippocampus and regions in the prefrontal cortex, including bilateral ventral prefrontal cortex, dorsal medial prefrontal cortex, and right lateral prefrontal cortex.

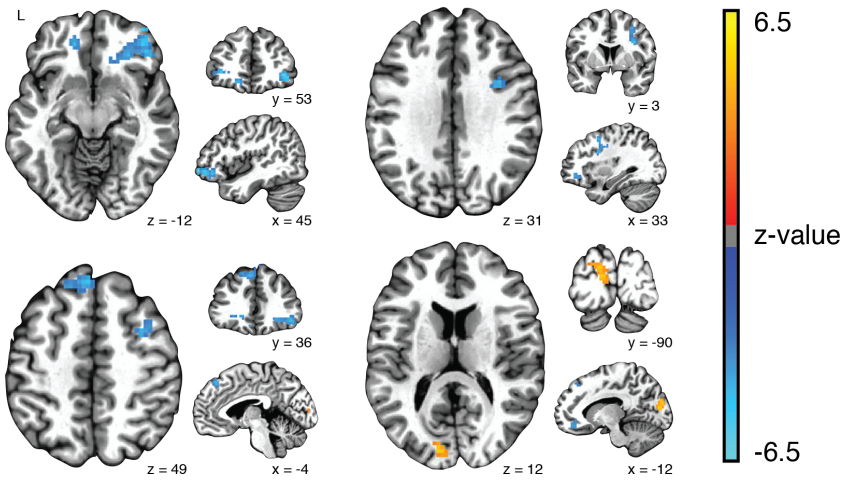


Figure 3.7. The relationship between performance improvement in the visuo-spatial task and hippocampal resting-state functional connectivity change during Visit 2. Change in functional connectivity between hippocampus and visual cortex was positively correlated with visuo-spatial learning. In contrast, functional connectivity change between hippocampus and several areas in prefrontal cortex was negatively correlated with visuo-spatial learning.

3.3.5.2 Motor sequence learning

To assess the impact of short-term motor-sequence training, I tested for an interaction between Visit 3 resting-state functional connectivity maps from AM and PM and motor-sequence learning performance, revealing a positive relationship between sensorimotor cortex resting-state functional connectivity increases and short-term motor-sequence learning in the right putamen as well as the right entorhinal cortex (Figure 3.8).

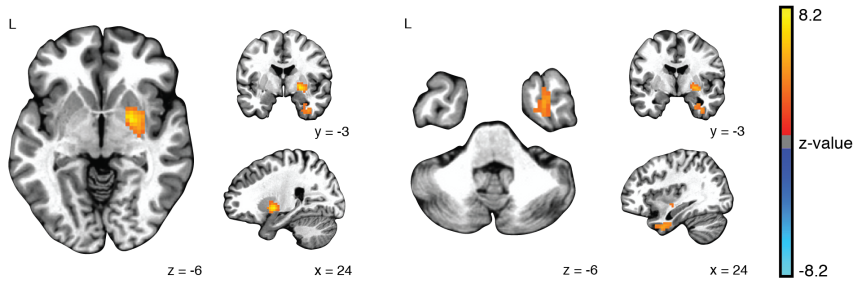


Figure 3.8. The relationship between short-term motor-sequence learning and sensorimotor cortex resting-state functional connectivity change during Visit 3. After motor-sequence training, an increase in resting-state functional connectivity between sensorimotor cortex and right putamen and entorhinal cortex was positively correlated with short-term motor-sequence learning.

3.3.6 Non-specific functional connectivity changes

Finally, I looked for regions that showed an apparent change in functional connectivity that did not depend on training. For this analysis, I looked for regions that showed a Main effect of Time-of-Day in the model that contained Visit and Time-of-Day as factors (Section 3.3.2).

Both hippocampus and sensorimotor cortices showed resting-state functional connectivity was impacted by Time-of-Day (Figure 3.9). Bilateral hippocampus showed an increase in resting-state functional connectivity from AM to PM with the middle cingulate cortex, superior temporal gyrus, occipital cortex, and right lateral prefrontal cortex. Similarly, resting-state functional connectivity

between sensorimotor cortex and bilateral superior temporal gyrus, occipital cortex, and a large region including inferior sensorimotor cortex and posterior insula, increased from AM to PM. However, resting-state functional connectivity between sensorimotor cortex and primary visual cortex and anterior cingulate cortex showed a decrease from AM to PM. These results suggest that measures of functional connectivity change systematically from AM to PM regardless of whether subjects are given training or not.

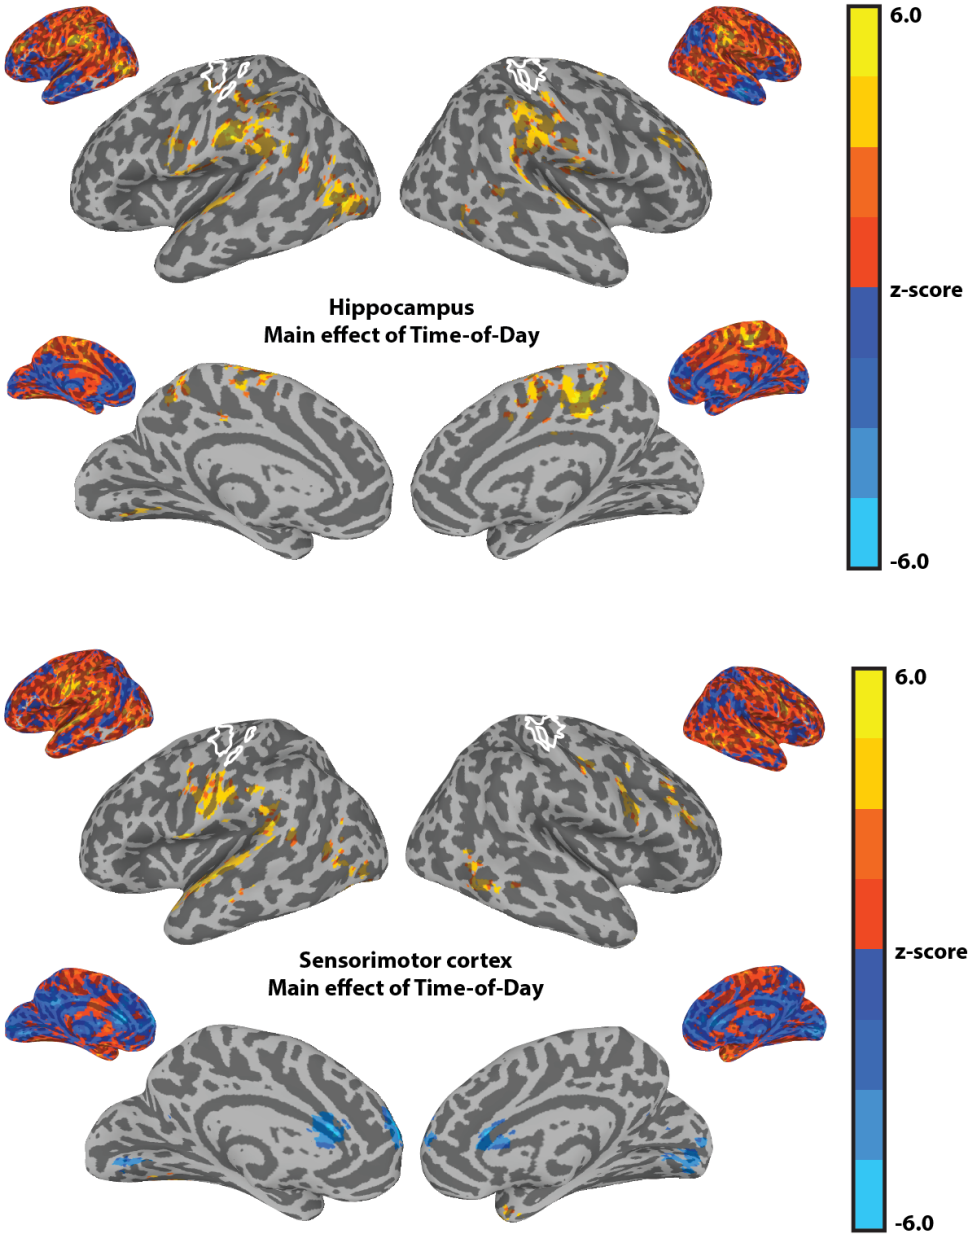


Figure 3.9. Non-specific changes in resting-state functional connectivity observed due to time-of-day for (upper) Hippocampus and (lower) Sensorimotor cortex.

3.4 Discussion

In this Chapter, my goal was to establish the sensitivity of resting-state functional connectivity to detect task-specific changes related to training. I found that by using resting-state functional connectivity measurements, I could detect changes specific to visuo-spatial and motor-sequence learning. In addition, I found that resting-state functional connectivity immediately after training reflected training-related benefits in behavioural performance. The importance of this result in the context of plasticity studies is discussed below.

3.4.1 Task-specific modulation after visuo-spatial training

After short-term training on the spatial layout of a race track, participants showed an increase in resting-state functional connectivity between the hippocampus and the posterior cingulate cortex. Recent studies have shown that posterior cingulate cortex plays a role in the representation of scenes [244-246], and that posterior parietal cortex shows task-related activity during a spatial navigation task similar to the one used here [247]. In addition, I observed increased resting-state functional connectivity between hippocampus and the putamen, and a decrease in resting-state functional connectivity with the sensory cortex and rostral dorsolateral prefrontal cortex. Given the role of the putamen in motor-sequence learning [248, 249], it may be that the increase in resting-state functional connectivity reflects cross-modal encoding between the spatial component of the task with the motor movements as the participant navigates the racetrack. Likewise, the concurrent decreases in resting-state functional connectivity may be due to prioritization of certain types of information during

memory consolidation [232]. I did not observe any relationship between the change in resting-state functional connectivity between these regions and change in behavioural performance, so I cannot say whether the observed resting-state functional connectivity change here is due to persistent rumination on the spatial layout of the track or whether offline memory processes were driving the interaction between hippocampus and PCC. Therefore, these findings should be interpreted with caution given the exploratory nature of this study.

Based on prior work that used the racing videogame paradigm [233, 247], I also looked for regions where resting-state functional connectivity with hippocampus correlated with visuo-spatial memory performance. I found a set of regions, including occipital cortex, ventral lateral prefrontal cortex, and medial prefrontal cortex where resting-state functional connectivity change after training was correlated with better memory for the spatial layout of the race track. This exploratory analysis suggests that prefrontal-hippocampal interaction may be associated with visuo-spatial memory, which is consistent with extant literature [233, 247]. These results provide motivation for future studies investigating the contribution of these functional connections to visuo-spatial memory in a more hypothesis-driven manner.

3.4.2 Sensorimotor cortex connectivity changes after motor-sequence learning

Prior literature suggests that training on an explicit motor sequencing task would evoke changes in resting-state functional connectivity between the basal ganglia and cerebellum [249]. The collective results from this study are in line with this hypothesis.

After training, functional connectivity between the sensorimotor cortex and caudate nucleus decreased after training. This decrease may reflect the diminished involvement of the caudate nucleus as performance approaches asymptote, which has been observed in both human and animal [250, 251]. To complement this finding, I found that a greater increase in connectivity between sensorimotor cortex and putamen was correlated to better performance after short-term training. Although these changes and relationships with behaviour are robust and specific, interpreting them is somewhat challenging. In the case of the caudate nucleus, I may be observing a change in resting-state functional connectivity that relates to behavioural performance prior to the functional connectivity measurement. With regard to the resting-state functional connectivity between sensorimotor cortex and putamen, it may be that training increases functional connectivity between these regions, and that sustained high resting-state functional connectivity in the offline period does correlate with performance after short-term training. However, the present data does not provide the means to test these hypotheses.

I also observed task-specific changes in resting-state functional connectivity between sensorimotor cortex and the lateral occipitotemporal cortex as well as left and right superior temporal gyri. However, the robustness of these changes is not readily apparent as they did not replicate in the LOSO analysis.

Similar to the findings regarding hippocampus resting-state functional connectivity and visuospatial performance, the regions that showed task-specific changes after motor-sequence learning did not show a significant correlation with behaviour. However, the whole-brain analysis of the relation between improvements in motors sequence performance and resting-state functional

connectivity change of the sensorimotor cortex during Visit 3 revealed a positive correlation in the right entorhinal cortex as well as the right putamen. As discussed above, the positive correlation in the putamen is in line with our expectation. With regard to the interpretation of the correlation between resting-state functional connectivity change and short-term motor-sequence learning, the medial temporal lobes (MTLs) (including entorhinal cortex) has been associated with explicit sequence learning [252, 253]. Together, the correlation analysis and the task-specific analysis, which revealed a reduction in resting-state functional connectivity between the caudate and sensorimotor cortex, provide a fuller picture of the likely role that components of the striatum play in motor learning.

3.4.3 Non-specific effects leading to functional connectivity change

In addition to training specific effects, I also observed non-specific changes in resting-state functional connectivity of both hippocampus and sensorimotor cortex that were present even on non-training visits. From morning to afternoon, resting-state functional connectivity of both hippocampus and sensorimotor cortex increased with temporal, parietal, insular and sensory cortices. Additionally, functional connectivity between sensorimotor cortex and the occipital and cingulate cortices decreased from morning to afternoon.

Similar time-of-day related fluctuations in resting-state functional connectivity have been reported using other analysis approaches, mostly focusing on changes in resting-state functional connectivity of regions associated with the default mode network. For example, one study found that resting-state functional connectivity between subnetworks within the default mode network tended to be unstable over the course of the day [254]. Consistent with this finding, another

study found resting-state functional connectivity was greater between the left and right MTL during the morning, but that resting-state functional connectivity between the MTL and other regions in the neocortex significantly increased in the afternoon [232]. This pattern of variation in resting-state functional connectivity from morning to afternoon was interpreted as a brain mechanism for consolidation of memory within a day. Another study used a seed-based correlation analysis and found a decrease in resting-state functional connectivity from morning to afternoon in cortical regions associated with the DMN [255]. This reduction in functional connectivity was related to a decrease in regional cerebral blood flow (rCBF) measured with arterial spin labelling. Outside the default mode network, one study found that the coherence and amplitude of low frequency fluctuations in resting-state fMRI varied significantly from morning to afternoon in widespread brain regions [256]. Thus, multiple studies have reported that resting-state functional connectivity fluctuates diurnally. The findings presented here extend these studies by showing that the time-of-day related change is not specific to regions associated with episodic memory. If left unaccounted for, these non-specific changes may be misinterpreted as changes due to training.

An important question that remains to be addressed concerns the mechanisms driving these non-specific changes. This study was not designed to elucidate the such mechanisms. However, the additional MRI data that was acquired during the same scan sessions as the resting-state fMRI data provide some insights [229, 257]. Briefly, analysis of the T_1-weighted images revealed that computed measures of grey matter volume and apparent cortical thickness showed a significant reduction in specific brain regions from morning to afternoon [229]. Analysis of the diffusion weighted images showed that the apparent

reduction in T_1-weighted measures is related to an increased diffusivity of a fast-diffusing pool of water that may include contributions not only from CSF, but also blood flow [257]. Interestingly, the spatial topography of the time-of-day related changes that were observed in T_1-weighted and diffusion-weighted images were most prominent in regions traversed by the leptomeningeal arteries, such as the interhemispheric fissure and the sylvian fissure. Qualitatively, the time-of-day related changes in the other MRI modalities and the resting-state functional connectivity changes reported here show a striking overlap, suggesting a shared physiological origin.

These observations, as well as the findings from other MRI studies of diurnal changes that have identified links between diurnal changes in regional cerebral blood flow and salivary cortisol [255], suggest that the diurnal fluctuations observed in the present study are likely to be driven by multiple physiological factors that include CSF and cerebral blood flow dynamics and mediated by neuroendocrine factors [255, 258]. Therefore, studies designed to tease apart the role of each of these factors are crucial to improve current understanding of the mechanisms underlying the phenomenon of diurnal fluctuations in MRI measures.

3.4.4 Limitations

One key limitation in the present study was the relatively low sample size (n=20), especially given the missing data. This was necessitated by the number of scans required and the length of the experimental sessions. Using a mixed-effects model to handle the missing data helped to address this issue. However, because of the low number of participants, I have interpreted the analyses of individual differences with particularly caution.

Another limitation is that, by design, the two tasks used here had substantially different demands. As such, I cannot speculate on the extent to which resting-state fMRI is sensitive to task-specific changes when the task demands are subtler. On this point, using lap-time improvement as a behavioural metric for the racing videogame may have been more a straight-forward means to compare the behaviour in the visuo-spatial task to the motor-sequencing task. However, lap-time could not be used because of the scene arrangement task: after the initial lap, participants may have altered their behaviour because they were aware of the scene test and focused on examining the course rather than driving quickly. The scene-arrangement task was chosen to be consistent with previous publications using this paradigm [233]. Subsequent studies might consider using lap-time as a metric for behavioural improvement in racing videogame tasks.

Finally, as mentioned above (Section 3.4.2), I found that the brain regions showing task-specific changes in resting-state functional connectivity at the group level did not show a correlation between functional connectivity change and individual differences in behavioural performance. The factors that lead to finding relationships between resting-state functional connectivity changes and changes in behavioural are complex, and caution is necessary when concluding from such results. I discuss this issue in more detail in Chapter 8 of this thesis.

3.5 Conclusions

In conclusion, in this Chapter I demonstrated that resting-state functional connectivity analyses may be sensitive to task-specific changes in functional connectivity induced by training. Interpreting changes in resting-state functional connectivity is not straightforward, given the prevalence of non-specific effects

that arise from the passage of time, as well as the effects of task performance that may not be relevant to learning, such as rumination [259]. These issues inform the discussion in Chapter 6, in which resting-state fMRI is used to investigate the effect of valenced feedback on functional plasticity after skill learning.

4 The impact of reward and punishment depends on task demands

Skills are learned to optimize particular behaviours given the constraints of the environment, which often includes rewards and punishments. However, the impact of exogenous feedback in the context of motor skill learning is not well characterized. In the present study, I investigated the effect of reward and punishment during training on both the serial reaction time task and the force tracking task. I found a dissociation of the effects of feedback on the tasks, primarily reflecting the impact of punishment. While punishment improved serial reaction time task performance, punishment impaired force tracking task performance. In contrast to prior literature, neither reward nor punishment benefitted memory retention. These results suggest that punishment impacts performance during training more than reward, and does so in a complex, task-dependent fashion.

4.1 Introduction

Reward and punishment, including biological reinforcers such as food, water, and pain, are important motivators for both human and animal behaviour. The majority of neuroscience research has focused on studying the effects of reward and punishment on decision-making [76, 260-262]. However, interest in using reward and punishment to augment motor learning has recently surged [17,

151, 152, 263-267]. Although motor sequence learning and motor adaptation learning are generally considered separate processes, the literature examining the impact of reward and punishment may be mutually informative. In the context of adaptation, it has been shown that punishment improves learning rate, while reward improves retention [151, 268] and also prevents anterograde interference from subsequent adaptations [269]. The benefit of reward for memory retention has also been shown for treadmill walking [265], as well. Perhaps most importantly, reward has been shown to renormalize adaptation learning in patients with cerebellar degeneration [267] and in stroke patients [266], which has led to optimism about the use of reward-augmented therapeutic regimes in clinical settings.

The effect of valenced feedback on motor skill learning has been less-well studied: one study reported that the benefit of punishment to performance and reward to immediate retention after SRTT learning [152], and another study reported that reward enhanced offline-gains and 30-day retention of a visuomotor tracking skill [17]. These two results are often cited together to support the role of reward in memory formation, and generally support the "simple heuristic" to explain the effect of reward versus punishment: that punishment improves performance while reward benefits memory. Citing the two studies together is misleading, however, as the studies by Wachter and colleagues [152] and Abe and colleagues [17] differ on several critical aspects. The main differences between the studies by Abe and colleagues and Wachter and colleagues are informativeness of feedback, retention test timing, the measure of sequence knowledge used.

Firstly, the informativeness of the feedback with respect to what was being trained differed between the two experiments. Wachter and colleagues used the

SRTT and provided feedback after each button press to inform the participant whether they had responded correctly and faster than their average RT during a pre-training familiarization period. Thus, in this study, participants could use the feedback instances to adapt their behaviour (accuracy or speed) at each discrete point in the 12-item sequence. In contrast, Abe and colleagues used the FTT and provided feedback at the end of the 9-second trial to inform the participant of their cumulative performance over the whole sequence. This does not actually provide information about when in the sequence of the movement the error occurred and therefore does not enable updating of performance.

Secondly, the timing of the retention tests differed substantially between the two studies. Abe and colleagues tested for delayed memory after immediately after learning, and again 6-hours, 24-hours, and 30-days after learning. By comparison, Wachter and colleagues only evaluated memory during an immediate retention test. Therefore, 'retention' is appreciably different across the two studies. For studies interested in a mechanistic explanation of the benefit of reward (e.g., improving acquisition or preventing forgetting), a consistent definition of 'retention' is quite important.

Finally, to assess sequence knowledge, Wachter and colleagues calculated the difference between RTs on the fixed-sequence and random-sequence trials, which is the canonical method for assessing learning on this task. Abe and colleagues instead used the improvement on tracking performance after training compared to their pre-training ability. Because no sequence-specific knowledge test was included in the study by Abe and colleagues, sequence-specific (or skill-knowledge from the SRTT perspective) cannot be dissociated from general improvement.

To address these concerns, in the present experiment, I investigated the effect of reward and punishment on these two tasks (the SRTT and FTT) with adapted implementations to facilitate comparison. In an initial training session, participants trained on either the SRTT or FTT and received feedback (monetary reward, monetary punishment, or motivated control) based on their performance (calculated as mean response time and accuracy for SRTT; mean distance from the target per block in the FTT). In both tasks, probe trials, during which stimuli were presented in either a fixed or a random order, were used to assess sequence-specific knowledge before and after training. I assessed general skill learning by comparing initial performance to performance after training, regardless of the probe block type. Skill retention was then probed in the absence of feedback at 1 hour, 24-48 hours, and 30 days after completion of the training.

Participants were able to learn both tasks successfully and the skill was almost entirely retained when tested after 30 days. Overall, I observed little effect of reward on either learning or retention. Punishment had no effect on skill retention, but had significant, task-dependent effect on learning. In the SRTT, training with punishment improved speed with minimal impact on accuracy. In contrast, training with punishment impaired performance on the FTT. These results suggest that the effect of feedback varies depending on the skill being learned, and while feedback impacts online performance, the benefit of reward reported to retention may be less robust than previously demonstrated.

4.2 Methods

The study design was the same for both tasks (Figure 4.1a-b). Participants trained on either the SRTT, or the FTT. For both tasks, trials were presented over

15 blocks. A 30-second break (minimum) separated each block of trials. Unbeknownst to the participants, during some blocks ("fixed sequence blocks") the stimulus would appear according to a repeating pattern (described below for each task). During other periods the appearance of the stimulus was randomly determined ("random sequence blocks").

Familiarization and training blocks were conducted in the bore of an MRI scanner. To acclimatize participants to the task, and establish their baseline level of performance, the task began with three random-sequence blocks without feedback ("familiarization blocks"). Participants were unaware of the forthcoming feedback manipulation during these familiarization blocks. Then the feedback period began, starting with a pre-training probe (three blocks, random – fixed – random), then the training blocks (six consecutive fixed-sequence blocks), and, finally, a post-training probe (three blocks, random – fixed – random). The difference in performance between the average of the two random blocks, versus the fixed sequence block, during the probes was used to index sequence knowledge [157].

To test the impact of reward and punishment on skill learning, participants were randomised into one of 3 feedback groups: reward, punishment, or uninformative (control). During the feedback period, reward, punishment, or control feedback was provided based on the participant's ongoing performance. The feedback paradigm for each task is outlined separately below.

Participants were given retention probes at 1 hour, 24-48 hours, and 30 days after training. No feedback was delivered during the retention probes. The second probe always occurred after at least one night's sleep.

The initial visit (Familiarization, Early probe, Learning, and Late Probe) took place while participants underwent MRI scanning (see Chapter 6).

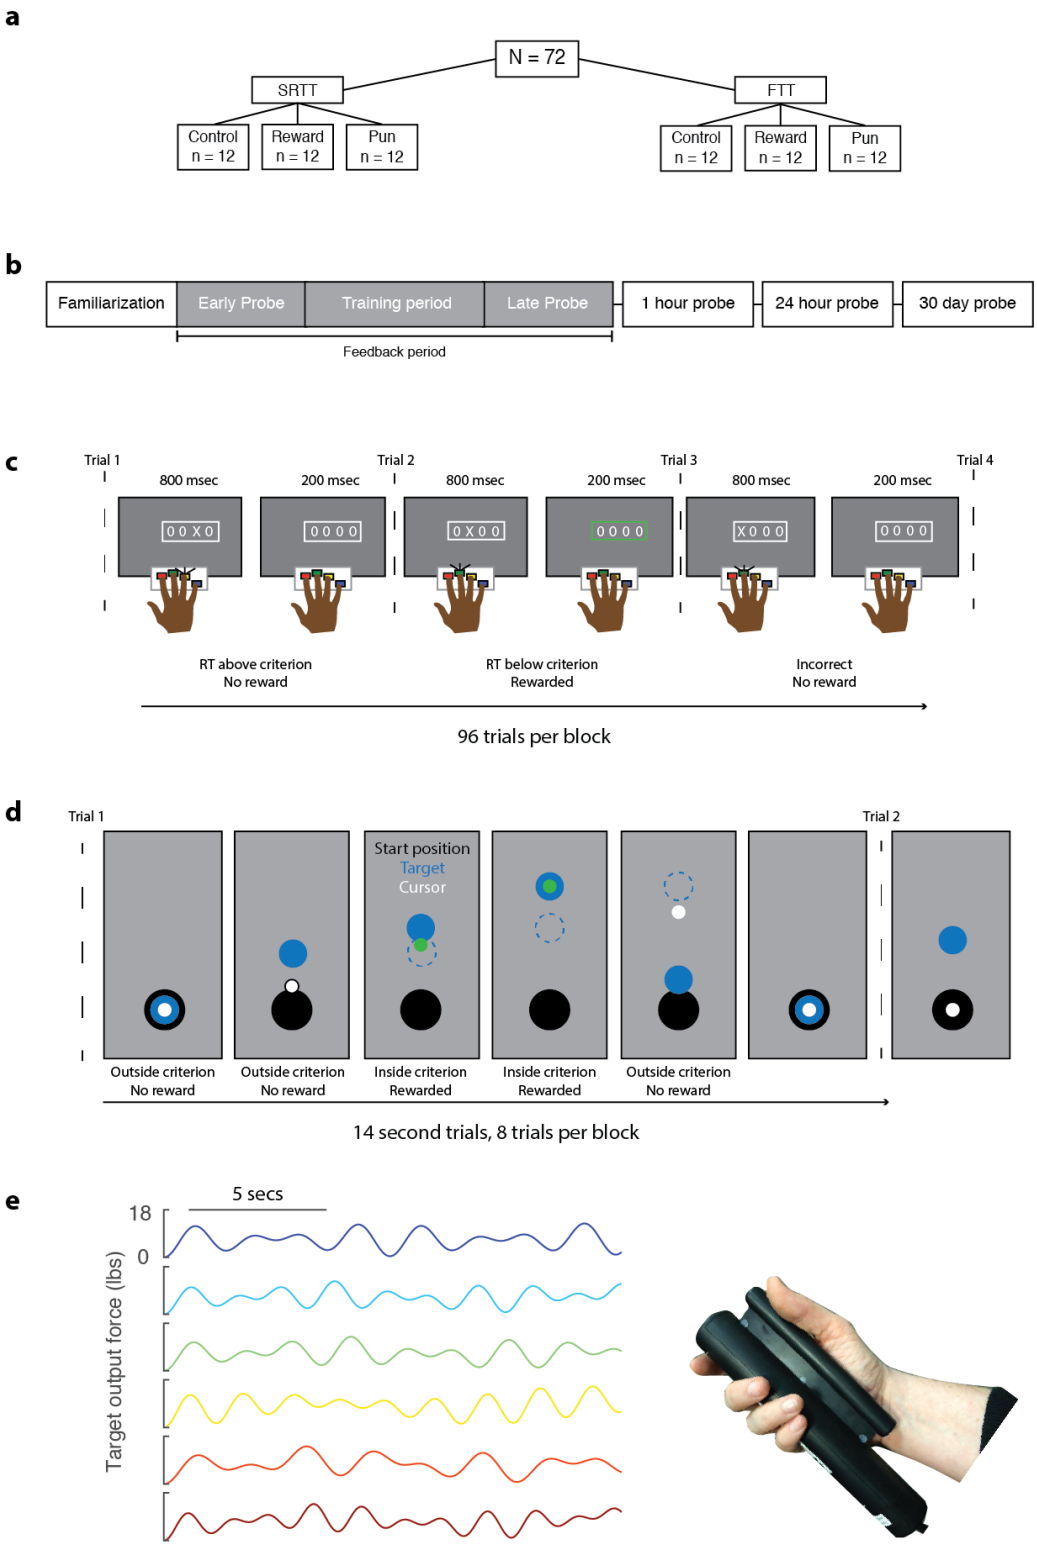


Figure 4.1. Experimental design. A. Seventy-two participants were divided between two skill learning tasks: the SRTT and the FTT. Within each task, participants were randomly assigned to three different feedback groups (reward, punishment, control). B. Experimental timeline. For each task, trials were grouped into blocks. Unbeknownst to the participants, during some blocks ("fixed-sequence blocks") the stimulus would appear according to a repeating pattern. During other periods the appearance of the stimulus was randomly determined ("random-sequence blocks"). Following familiarization blocks, participants were trained on the task with feedback. To assess sequence knowledge, training was bookended by early and late probes in which participants performed three blocks arranged random - sequence - random. Participants were then tested for sequence knowledge without feedback 1 hour, 24 hours, and 30 days after learning. C. SRTT. Participants were presented with four locations on a screen denoted by "O's". A trial began when one "O" changed to an "X". Participants were instructed to press the corresponding button on a controller as fast and accurately as possible. After 800 ms, the X changed back to an O, and participants were given feedback for their performance on that trial. Performance in the SRTT was based on RT and accuracy. If participants responded accurately and were faster than their criterion RT, a participant would receive positive feedback (reward, or absence of punishment) on that trial. If they were slower or inaccurate, they would receive the negative outcome (either punishment or absence of reward). D. FTT. Participants held a force transducer in their right hand and saw a black circle (start position), a blue circle (target), and a white circle (cursor). Participants were instructed to squeeze the force transducer to keep the cursor as close to the center of the target as possible. The target moved continuously during the trial (12 seconds), followed by a 2 second break between trials. The distance of the cursor from the target was the measure of performance. If the participant was closer to the center of the target than they were on their previous block of trials, they would receive positive feedback. During sequence blocks the target followed one of six trajectories, (D, left) whereas during random blocks the target would follow a random trajectory.

4.2.1 Participants

78 participants (47 female, mean age = 25 years ± 4.25 standard deviation) participated in this experiment. All participants were right-handed, free from neurological disorders, and had normal or corrected-to-normal vision. All participants gave informed consent and the study was performed with National Institutes of Health Institutional Review Board approval in accordance with the Declaration of Helsinki (93-M-0170, NCT00001360). Data from six individuals were removed from the study due to inattention (defined as non-responsive or

inaccurate on greater than 50% of trials during training) or inability to complete the training session. Participants were compensated using the local currency ($).

4.2.2 Serial reaction time task

The version of the SRTT used here adds feedback to the traditional implementation. At the beginning of each block participants were presented with four "O"s, arranged in a line, at the centre of the screen. These stimuli were presented in white on a grey background (Figure 4.1c). A trial began when one of the "O"s changed to an "X". Participants were instructed to respond as quickly and accurately as possible, using the corresponding button, on a four-button response device held in their right hand. The "X" remained on screen for 800 ms regardless of whether the participant made a response, followed by a 200 ms fixed inter-trial interval, during which time the four "O"s were displayed.

A block consisted of 96 trials. During fixed-sequence blocks, the stimuli appeared according to a fixed 12-item sequence repeated 8 times (e.g., 3-4-1-2-3-1-4-3-2-4-2-1). Each fixed block began at a unique position within the sequence, to help prevent explicit knowledge of the sequence from developing [252]. In the random blocks, the stimuli appeared according to a randomly generated sequence, without repeats on back-to-back trials, so, for example, participants would never see the triplet 1-1-2.

Between each block, participants saw the phrase "Nice job, take a breather". After five seconds, a black fixation-cross appeared on the screen for 25 seconds. Five seconds before the next block began, the cross turned blue to cue the participants that the block was about to start.

During the retention probes, participants performed three blocks (random – fixed – random on a 15-inch MacBook Pro using a button box identical to the one

used during training. During these retention probes, the next trial began 200 ms after the participant initiated their response. No feedback was given during the retention blocks. The first button press made after stimulus presentation was considered the participant's response. Any missed trial was counted as an error, and only correct trials were considered for analysis of RTs.

4.2.3 Force tracking task

In the FTT, participants continuously modulated their grip force to match a target force output [15, 16]. In the traditional implementation, participants are exposed to a single pattern of force modulation repeated each trial. This design does not allow discrimination between general improvement (i.e., familiarization with the task and/or the force transducer) and improvement specific to the trained sequence of force modulation. Therefore, I decided to adapt the traditional FTT method to align it with the experimental design that is traditional for the SRTT, i.e., by including random sequence blocks.

A given trial consisted of a 14 second continuous pattern of grip modulation. At the beginning of a trial, participants were presented with three circles on a grey background projected onto a screen: a white circle (Cursor, 0.5 cm diameter), a blue circle (Target, 1.5 cm diameter), and a black circle (bottom of the screen, 2 cm diameter, indicating the position corresponding to minimum pressure; Figure 4.1d). Participants held the force transducer (Current Designs, Inc., Philadelphia, PA) in the right hand between the four fingers and palm (Figure 1e, right). Participants were instructed to squeeze the force transducer (increasing force moving the cursor upwards) to keep the cursor as close to the centre of the target as possible as the target moved vertically on the screen. During fixed blocks, participants were randomly assigned to one of six sequences (Figure

4.1e, left). During random blocks, the target followed a trajectory generated by the linear combination of four waveforms, with periods between 0.01 and 3 Hz. These waveforms were constrained to have identical average amplitude (target height), and the number and value of local maxima and minima are consistent with the sequence blocks.

For data analysis, the squared distance from the cursor to the target was calculated at each frame refresh (60 Hz). The first 10 frames were removed from each trial. The mean of the remaining time points was calculated to determine performance, and trials were averaged across blocks.

4.2.4 Feedback

Because this study took place in the MRI scanner, all participants were paid a base remuneration of $80 for participating in the study. At the start of the feedback period, participants were informed they could earn more money based on their performance.

During the feedback period, participants were given either reward, punishment, or control feedback. The presence of reward or the absence of punishment was based on a participant's performance. In both the SRTT and the FTT, the initial criterion was based on a participant's median performance during the final familiarization block. To encourage continuous improvement, this criterion was re-evaluated after each block. In the reward group, the feedback indicated that the participant's performance was improving. In the punishment group, the feedback indicated they were getting worse. Because the frequency of feedback events differed between the reward and punishment groups (reward from high-to-low as training progressed, punishment from low-to-high), the control group was split into two different sub-groups (control-reward and

control-punishment). The control groups received feedback at a frequency that matched the corresponding feedback group but was not related to their performance. Participants in the control group were made aware that the feedback was not meaningful. I considered the reward and punishment control groups together in the analyses, as is typical in these studies [17, 152].

4.2.4.1 SRTT feedback

In the SRTT, performance was defined as the accuracy (correct or incorrect) and RT of a given trial. Feedback was given on a trial-by-trial basis (Figure 4.1c). This was indicated to the participant when the white frame around the stimulus changed to green (reward) or red (punishment). In the reward group, the participants were given feedback if their response was accurate and their RT was faster than their criterion RT, which indicated that they earned money ($0.05 from a starting point of $0) on that trial. In the punishment group, participants were given feedback if they were incorrect or their RT was slower than their criterion, which indicated that they lost money ($0.05 deducted from a starting point of $55) on that trial. Participants in the control-reward and control-punishment groups saw red or green colour changes, respectively, at a frequency matched to the punishment and reward groups. Control participants were told that they would be paid based on their speed and accuracy. Importantly, to control for the motivational differences between gain and loss, participants were not told the precise value of a given trial. This allowed us to assess the hedonic value of the feedback, rather than the level on a perceived-value function. Between blocks, for the reward and punishment groups, the current earning total was displayed (e.g., "You have earned $5.00"). Control participants saw the phrase, "You have earned

money." The criterion RT was calculated as median performance in the first familiarization block. After each block, the median + standard deviation of performance was calculated, and compared with the criterion. If this test criterion was faster the previous criterion, the criterion was updated. Only the correct responses were considered when establishing the criterion RT.

4.2.4.2 FTT feedback

Feedback in the FTT was based on the distance of the cursor from the target (Figure 4.1d). For the reward group, participants began with $0. As participants performed the task, their cursor turned from white to green when the distance from the target was less than their criterion. This indicated that they were gaining money at that time. In the punishment group, participants began with $45, and the cursor turned red if it was outside their criterion distance. This indicated that they were losing money. For reward-control and punishment-control, the cursor changed to green or red, respectively, but was unrelated to their performance. For control, the duration of each feedback instance, as well as cumulative feedback given on each trial, was matched to the appropriate group. Between each block, participants were shown their cumulative earnings. Control participants saw the phrase "You have money." The criterion distance from the target was calculated as median performance in the first familiarization block. After each block, the median + standard deviation of performance was calculated, and compared with the criterion. If this test criterion was more accurate (FTT) than the previous criterion, the criterion was updated.

4.2.5 Statistical analysis

In both tasks, the six training blocks were compared using a repeated-measures ANOVA to establish differences in learning rate (Block x Group). Learning was indexed by comparing the performance (RT and accuracy separately for SRTT; squared distance from the target [squared error] for FTT) on the fixed-sequence blocks to the average of the two random-sequence blocks at the pre- and post-training time points using a repeated-measures ANOVA (Time point x Sequence x Group). Sequence-specific memory was evaluated by comparing the fixed-sequence block, to the mean of the two random blocks, at each retention time point using a repeated-measures ANOVA (Time point x Sequence x Group). A Bonferroni correction was applied for post-hoc analyses to correct for multiple comparisons. If sphericity was violated, the Hyunh-Feldt correction was applied.

4.3 Results

4.3.1 Punishment improves online series reaction time task performance

Data from the SRTT are presented in Figure 4.2. I investigated the impact of reward and punishment on SRTT sequence learning in three different ways. First, I compared sequence knowledge during sequence knowledge probes either early in learning (immediately following familiarization when feedback was first introduced) or late in learning (at the end of the training session). During these probes, I estimated sequence knowledge by calculating RT difference between fixed- and random-sequence blocks. A repeated measures ANOVA, with Group (reward, punishment, control), Sequence (fixed, random) and Time-point (early, late) as factors revealed a significant three-way interaction between Group,

Sequence, and Time-point ($F_{(2,33)} = 5.370$, $p < 0.01$). Follow-up analyses indicated that both punishment and reward groups acquired more skill knowledge during the early sequence knowledge probe than control ($F_{(2,33)} = 5.213$, $p < 0.05$; punishment v control: $t_{(22)} = 3.455$, $p < 0.005$, reward v control: $t_{(22)} = 2.545$, $p < 0.02$), but did not differ from each other (reward v punishment: $t_{(22)} = 0.707$, $p = 0.487$). Further, the control group evidenced a greater gain in sequence knowledge from the early to the late sequence knowledge probe compared to reward ($t_{(22)} = 2.884$, $p < 0.01$), although this comparison was not significant for control versus punishment when correcting for multiple comparisons ($t_{(22)} = 2.075$, $p = 0.05$). This may partially reflect the benefit of valenced feedback to early learning. These results suggest that feedback facilitates rapid sequence learning on the SRTT. Notably, regardless of feedback, all participants were faster in both the fixed and random sequence blocks post-training, compared to pre-training, (Main effect of Time-point: $F_{(1,33)} = 30.903$, $p < 0.001$); as expected, participants were faster during fixed- than random-sequence blocks (Main effect Sequence: $F_{(1,33)} = 58.412$, $p < 0.001$). Furthermore, sequence-specific knowledge increased over time (Time-point x Sequence interaction: $F_{(2,33)} = 9.291$, $p < 0.005$), as the difference in RT between fixed- and random-sequence blocks increased from the pre- to post-training probe ($t_{(35)} = 2.815$, $p < 0.01$).

Second, to examine the effect of feedback on learning rate, I compared the median RT across the six consecutive sequence training blocks immediately following the early sequence knowledge probe using a repeated measures ANOVA with Block (1-6) and Group as factors. All participants showed improvement over the course of training (Main effect of Block: $F_{(5,33)} = 11.224$, $p < 0.001$). I also found a Main effect of Group ($F_{(2,33)} = 3.286$, $p < 0.05$) and follow-up tests indicated that

the punishment group was significantly faster than the control group overall (punishment versus control: $t_{(22)} = 2.884$, $p < 0.012$), but there was no difference between the reward and control groups ($t_{(22)} = 0.480$, $p = 0.636$), or the reward and punishment groups ($t_{(22)} = 1.757$, $p = 0.093$, two-tailed) during the training period. The lack of a significant Group x Sequence interaction in the post-training probe highlights that this is a general, rather than sequence-specific, improvement.

Finally, I examined the impact of feedback on retention. All groups demonstrated retention of sequence knowledge at all time-points (Main effect Sequence: $F_{(1,33)} = 100.245$, $p < 0.001$; $t_{(35)} = 10.036$, $p < 0.001$). There was no influence of feedback group on retention.

Collectively, these results show that both reward and punishment increased early learning of the sequence with punishment additionally having a marked effect on performance during training.

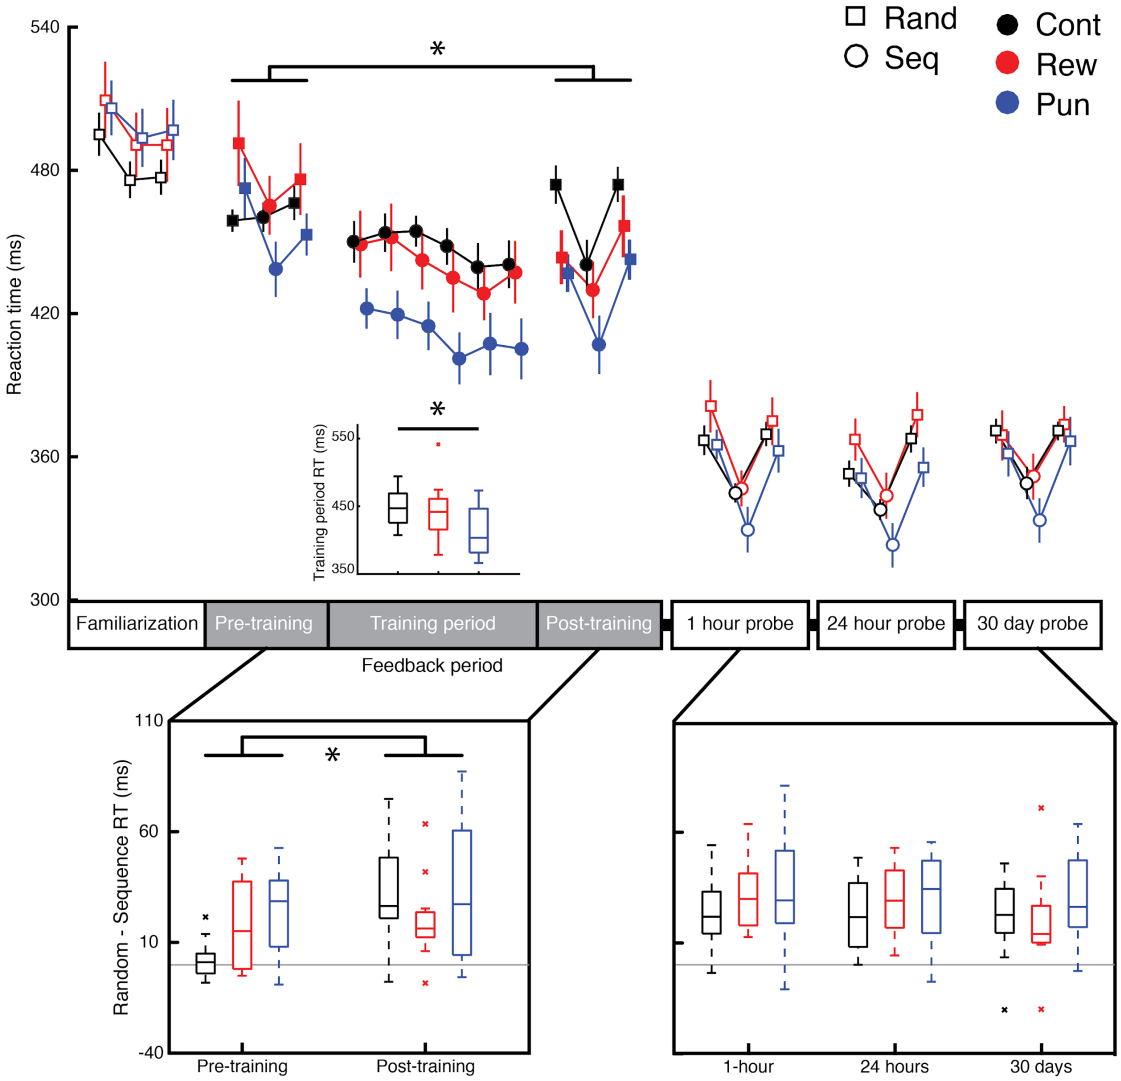


Figure 4.2. Punishment improves performance of the SRTT. During the training period, the punishment group was significantly faster than the control group overall. There was no difference between the reward and control groups or the reward and punishment groups during the training period. The reward group showed a greater reduction in average RT from pre- to post-training than the control group. Both the punishment and the reward groups acquired more skill knowledge during the early sequence knowledge probe than control but did not differ from each other. The punishment group did not differ from the control group when correcting for multiple comparisons. Feedback did not affect retention at any time-point (lower right panel). Main panel shows mean ± SEM. Box plots show median, crosses show within group outliers. Asterisks denote periods with significant effects of feedback (p< 0.05).

4.3.1.1 Error-rates did not differ between the reward and punishment groups

Although RT is the primary measure of behaviour in the SRTT, for completeness I examined the effect of feedback on accuracy during training (Figure 4.3). Before feedback was introduced, there was no significant difference between the feedback conditions ($F_{(2,35)}$ = 0.795, p = n.s.). I first compared the impact of feedback on the sequence knowledge after learning by comparing the difference in the proportion of correct trials during the pre- and post-training probes using a repeated measures ANOVA, with Group (reward, punishment, control), Sequence (fixed, random), and Time-point (pre, post) as factors. As expected, participants made less errors in fixed compared to random blocks (Main effect Sequence: $F_{(1,33)}$ = 40.474, p < 0.001). Participants were less accurate overall during the post-training Time-point (Main effect of Time-point: $F_{(1,33)}$ = 37.832, p < 0.001).

The type of feedback given during the probe blocks did affect the accuracy during these blocks (Main effect of Group: $F_{(1,33)}$ = 5.335, p < 0.01). Post-hoc tests revealed that the control group was significantly more accurate than the reward group ($t_{(22)}$ = 2.743, p < 0.012) and the punishment group ($t_{(22)}$ = 3.5, p < 0.005), while the reward and punishment groups did not differ ($t_{(22)}$ = 0.362, p = 3.720.). Group showed a significant interaction with Time-point ($F_{(2,33)}$ = 11.894, p < 0.001). Post-hoc analyses revealed that the reward and punishment groups made more errors during the post-training probe compared to the pre-training probe, indicating that the sacrificed accuracy more than the control group (Control v Reward: $t_{(22)}$ = 2.658, p < 0.014, Control v Punishment: $t_{(22)}$ = 4.562, p < 0.001). The change in the proportion of correct trials from pre- to post-training was not

significantly different between the punishment and reward groups when correcting for multiple comparisons ($t_{(22)}$ = 2.547, $p_{uncorrected}$ = 0.018). There was no significant interaction between Group and Sequence ($F_{(2,33)}$ = 2.548, p = 0.093) or for Group x Sequence x Time-point ($F_{(2,33)}$= 1.659 p = 0.206).

As I did for RT, I examined the effect of feedback on the learning period by comparing accuracy across training period using a repeated measures ANOVA with Block and Group as factors. Participants tended to get less accurate over time (Main effect Block: $F_{(5,165)}$ = 5.078, p < 0.001; Block 1 v Block 6: $t_{(35)}$ = 3.209, p < 0.005). There was a Main effect of Group on accuracy during these blocks ($F_{(2,33)}$ = 10.768, p < 0.001). Post-hoc analysis revealed that the punishment group ($t_{(22)}$ = 5.462, p < 0.001) and the reward group ($t_{(22)}$ = 3.140, p < 0.005) were less accurate than the control group during the learning period. The reward and punishment groups did not differ in their accuracy ($t_{(22)}$ = 1.182, p = 0.251). There was also a significant Block by Group interaction ($F_{(10,33)}$ =3.419, p < 0.001). Follow-up analysis showed that the punishment and reward groups showed a greater decrease in accuracy from Block 1 to Block 6 compared to the control group (Reward v Control: $t_{(22)}$ = 2.852, p < 0.01; Punishment v Control: $t_{(22)}$ = 2.851, p < 0.05) but were not significantly different from each other (Reward v Punishment: $t_{(22)}$ = 0.0, p = 1).

Finally, to test whether the feedback groups differed during the delayed retention probes, I compared the average accuracy during the fixed-sequence test probes versus the accuracy during random-sequence test probes using a repeated measures ANOVA with Probe, Group, and Sequence as factors. All subjects showed retention of sequence knowledge during the delayed probes, evidenced by elevated accuracy in fixed- compared to random-sequence blocks (Main effect of

Sequence: $F_{(1,33)} = 24.492$, $p < 0.001$; $t_{(35)} = -4.954$, $p < 0.001$). There was also a main effect of Probe ($F_{(2,33)} = 6.814$, $p < 0.005$), indicating that participants tended to be more accurate at the 30 day test compared to the 1 hour test ($t_{(35)} = -2.554$, $p < 0.05$). There was no effect of Feedback Group on accuracy in the delayed test probes.

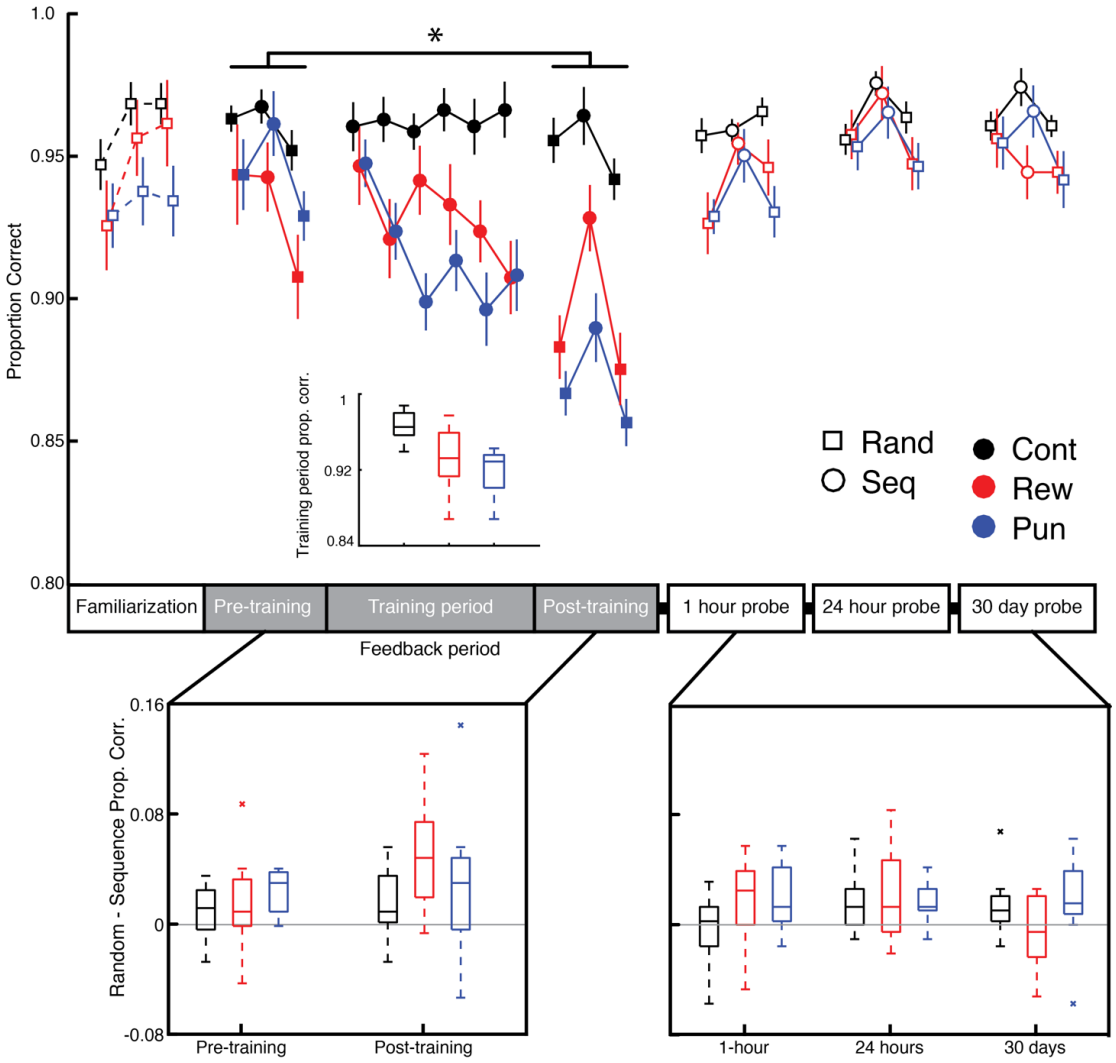


Figure 4.3. Error-rate during the SRTT. Although the primary measure of behavior in the SRTT, for completeness I examined the accuracy during training. Before feedback was introduced, there was no significant difference between the feedback conditions. As expected, participants made less errors in fixed- compared to random-sequence blocks. Participants were less accurate overall during the post-training time-point. The type of feedback given during the probe blocks did affect the accuracy during these probe blocks. Post-hoc tests revealed that the control group was significantly more accurate than the reward and punishment groups, while the reward and punishment groups did not differ. The reward and punishment groups made more errors during the post-training probe compared to the pre-training probe, indicating that they sacrificed accuracy more than the control group. Feedback did not affect accuracy during the delayed test probes.

4.3.2 Punishment impairs force tracking task performance

I conducted the same three analyses as detailed above for the SRTT on data from FTT (Figure 4.4). First, sequence knowledge during the feedback period was evaluated by comparing the mean squared error during fixed-sequence and random-sequence blocks in the early and late sequence knowledge probes using repeated measures ANOVA (Time-point x Sequence x Group). I found a significant interaction between Time-point and Group interaction ($F_{(2,33)} = 3.526$, $p < 0.05$), but unlike the SRTT, there was no three-way interaction with Sequence ($F_{(2,33)} = 1.212$, $p = 0.310$). Follow-up analyses indicated that only the reward group improved from pre- to post-training time-point (pre- versus post-training probe, 1-sample t-test, Reward: $t_{(11)} = 4.250$, $p < 0.001$, Punishment: $t_{(11)} = 0.100$, $p = 0.922$, Control: $t_{(11)} = 2.292$, $p = 0.043$ [n.s. corrected for multiple comparisons]). The punishment group showed significantly less improvement than the reward group from the pre- to post-training probe ($t_{(35)} = 2.372$, $p < 0.03$) but was not significantly worse than the control group ($t_{(35)} = 1.206$, $p = 0.241$). There was no difference between the reward and control groups ($t_{(35)} = 1.688$, $p = 0.106$). Unlike the SRTT, there was no significant Time-point x Sequence x Group interaction, suggesting that feedback modulated general, rather than sequence-specific, learning on the FTT.

Second, I examined performance during the six consecutive sequence training blocks using repeated measures ANOVA, with Block and Group as factors. All feedback groups showed improvement across the training period (Main effect of Block: $F_{(5,165)} = 8.478$, $p < 0.001$; S2 versus S7: $t_{(35)} = 2.836$, $p < 0.01$). Although the reward group tended to outperform the punishment group during training,

there was no effect of Group on learning rate in the FTT (Group x Block: $F_{(10,165)}$ = 1.186, p = 0.156).

Finally, I examined the effect of feedback on retention in the FTT. Five participants did not complete the retention probes due to timetabling. This left us with 10 control, 9 reward, and 11 punishment participants for retention analyses. All groups demonstrated retention of sequence knowledge at all time-points (Main effect of Sequence, F (1,27) = 86.387, p < 0.001; $t_{(35)}$ = 9.030, p < 0.001). There was no main effect or interaction with feedback group on retention.

Collectively, these results show that the primary effect of feedback in FTT was for punishment to impair learning from the pre- to post-training probe time-points.

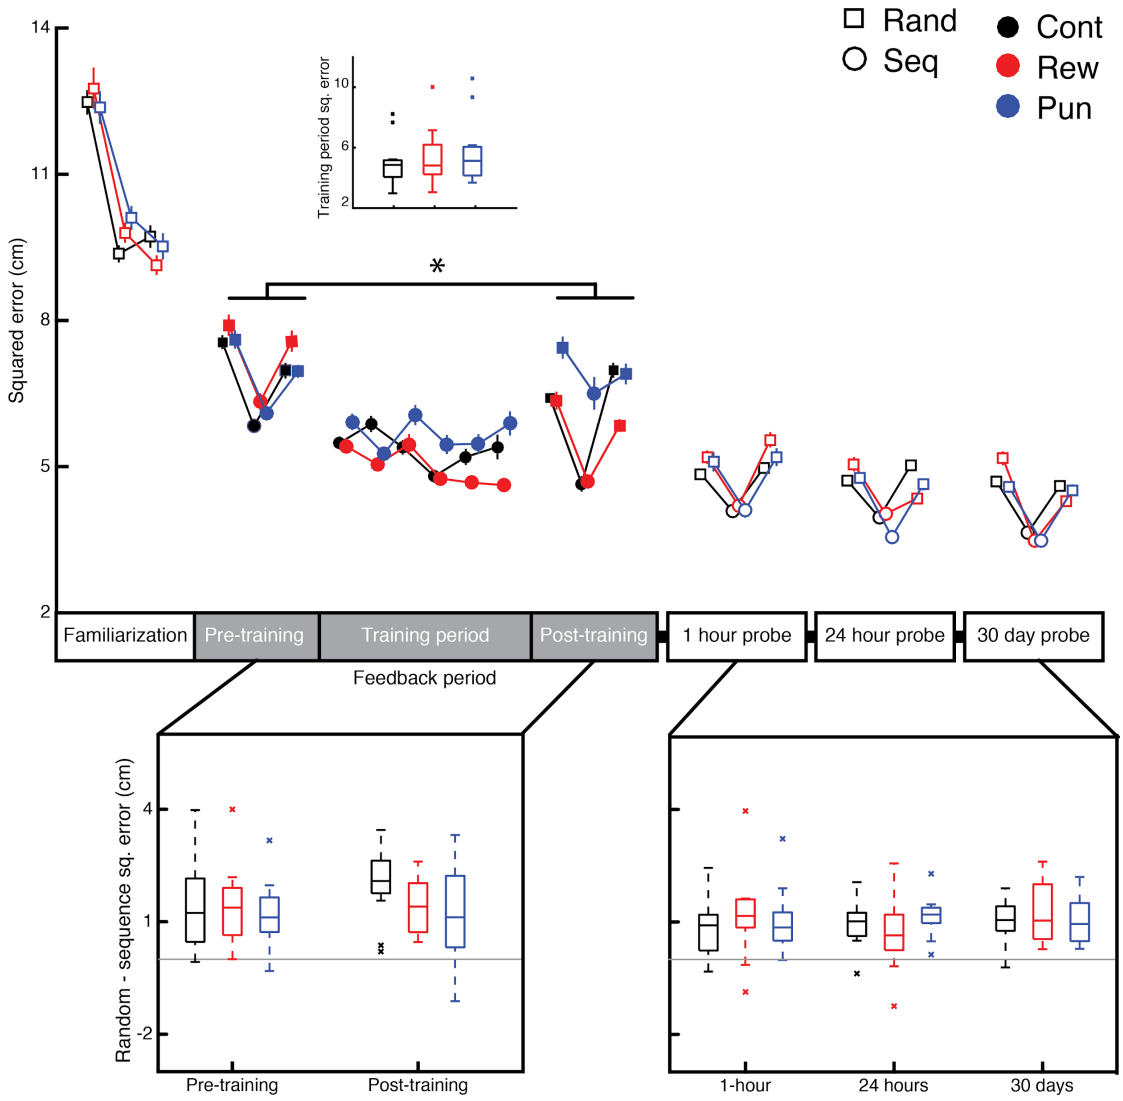


Figure 4.4. Punishment impairs to performance on the FTT. Compared to reward, punishment impaired general performance improvement from the pre- to post-training probe. Reward was not beneficial compared to control. Punishment did not impair performance compared to control. Unlike the SRTT, feedback had no influence on sequence specific knowledge (lower left) or performance during the training period (inset). As was found in the SRTT, feedback did not affect performance during the retention probes (lower right). Main panel shows mean ± SEM. Box plots show median, crosses show within group outliers. Asterisks denote periods with significant effects of feedback (p< 0.05). One outlier from punishment in the post-training period is not pictured in the boxplot [Random – Sequence = -7 cm sq. err].

4.3.3 Conscious recall on the serial reaction time task

I assessed conscious recall for the SRTT, but not the FTT. For the SRTT, because I did not reveal the presence of the sequence until 30 days after learning, I believed that verbal recall would be an insensitive measure to assess conscious

knowledge. Therefore, I implemented the process dissociation procedure in the same manner described by Song and Cohen (2015). Briefly, after completion of the 30 day test probe, I asked participants to input 100 button-presses while trying to include as much of the sequence as possible. They were specifically instructed that if they thought they thought that they knew any of the sequence, or several chunks of the sequence, that they should input those pieces as many times as they felt that they could. The area under the curve (AUC) for each number of sequential items (triplets, quadruplets, quintuplets through to dodecs [runs of 12]) included in the 100 button presses generated by the subject was compared to chance, which was estimated via 10,000 Monte-Carlo simulations of 100 non-repeating items. For these simulations, item frequency was calibrated to each subject using the frequency of each button's appearance during 100 button presses where the subject was instructed to press buttons randomly with no repeats. The groups were compared using a one-way ANOVA with Group as a factor.

Using this method, I determined that there was no difference in conscious recall between the feedback groups ($F_{(2,35)} = 0.705$, $p = 0.501$). The group mean AUC difference from chance was 0.08 ± 0.4 (mean ± standard deviation), and prior work has suggested that participants who are informed of the presence of a sequence prior to training have a mean AUC difference of approximately 4 [270], and therefore I believe that the participants in this study showed unconscious learning when assessed at 30 days after training.

4.3.4 Earnings

The amount of money earned in the reward and punishment groups did not differ in the SRTT ($t_{(22)} = 0.096$, $p = 0.924$), but did differ in number of feedback

instances ($t_{(22)}$ = -5.604, $p < 0.001$), as punishment was significantly better at avoiding negative feedback than reward was at obtaining positive feedback. Reward and punishment did not differ significantly in the amount of feedback received during FTT ($t_{(22)} = 0.275$, $p = 0.786$), but did differ in total amount earned (Reward = \$19.55 ± 0.45, Punishment = \$25.75 ± 1.08; $t_{(22)}$ = -5.684, $p < 0.001$).

4.4 Discussion

In this Chapter, I sought to determine whether the impact of reward and punishment generalizes across different types of motor skill tasks, such as the SRTT and FTT used here. I found that between the two skills, training with punishment had opposing effects on performance. During performance of the SRTT, training with punishment led to improved RTs overall with minimal detriment to accuracy compared to the reward group. In contrast, training with punishment impaired performance of the FTT. These effects were only present during the period with feedback; there was no effect of training with feedback on general or sequence-specific retention measured at 1 hour, 24 hours, and 30 days in either task. These results refute any simple model of the interaction between feedback and performance. Instead, these data suggest that the impact of valenced feedback, and punishment in particular, depends on the skill being learned.

There may be a number of reasons for this task-specific effect of feedback. While both tasks rely on sequence learning, they differ with respect to the mechanism that facilitates improvement. The motivational salience of punishment (i.e., loss aversion) may explain the performance benefit seen on the SRTT, where the added attention facilitated by punishment has been hypothesized to recruit additional neural resources to aid SRTT performance [152, 271].

However, a purely motivational account cannot explain the negative effect of punishment during FTT performance. Therefore, researchers examining the effect of reward and punishment need to consider alternative explanations that may account for the differential effects of reward and punishment to performance these two tasks.

One dimension on which the two tasks differ is in their motor demands. Specifically, in these implementations, performance on the FTT relies on more precise motor control than the SRTT, which requires only a simple button press. Within the motor system, others have reported that reward-related dopaminergic activity reduces motor noise [272], while dopaminergic activity associated with punishment leads to an increase in motor variability, i.e., noise [273]. Consistent with increased motor noise, I found that punishment impaired general (i.e., non-sequence-specific) performance on the FTT. During the retention test without feedback after one hour, the punishment group performed as well as the reward and control groups, suggesting that the increased motor noise may not impair learning, per se. Rather, the motor noise may impair general performance. Because increased motor variability was not directly measured in the SRTT, participants would not be penalized for any additional variation in movement that did not directly impact RT. If I had assessed motor variability during SRTT performance, I may have found that punishment induced more variation in movement on this task as well. Future work should examine this question.

The implementations of the tasks used here also differed with respect to the information content of a given instance of feedback. Learning on the SRTT is thought to rely on the positive prediction error encoded in the striatum that occurs on fixed-sequence trials [152, 274]. The reward or punishment in the SRTT

may augment this positive prediction error and facilitate performance and learning. In contrast, the moment-to-moment feedback given on the FTT is not associated with an instantaneous positive prediction-error signal. Rather, this implementation of the FTT is similar to motor tasks that rely on the cerebellum and may therefore not benefit from moment-to-moment feedback [275]. Although information content was not intentionally manipulated, this difference may also alter effect the reward and punishment on these tasks.

Unlike prior studies, I saw no benefit of reward to retention [17, 151, 152, 276]. Most studies examining reward and punishment in skill learning have only considered immediate recall [151, 152, 276], and only one study has shown a benefit of reward to long-term retention of a motor skill [17]. In their study, Abe and colleagues observed that participants training with control and punishment feedback had diminished performance after 30 days compared to their post-training time-point (i.e., they evidenced forgetting). Importantly, the reward group in the study by Abe and colleagues study showed offline gains from the immediate time-point to 24 hours after training that lasted 30 days. In my study, the punishment and control group did not evidence forgetting from 24-48 hours to 30 days, the reward group did not show any offline-gains. As such, I am confident that feedback did not alter retention in my study.

While not discussed at length by Abe and colleagues, the punishment group in that study performed significantly worse during training, suggesting that the quality of acquisition may have differed across the groups. Therefore, it is unclear whether the difference in memory observed in their study can be attributed to a benefit of reward to consolidation or to ineffective acquisition when training with punishment. The present study design differed from the implementation used by

Abe and colleagues with respect to feedback timing and trial duration. However, these results question the robustness of the finding that reward benefits skill retention.

With respect to the SRTT, it is worth considering that the participants in my experiment evidenced less sequence-specific learning than some others have found in unrewarded versions of this task, where the difference between fixed-sequence and random-sequence trials can be up to 80 ms [277, 278]. However, there is considerable variability in the difference between fixed-sequence and random-sequence trials on the SRTT reported in the literature, and some groups have reported sequence-specific learning effects on the SRTT to be between 10 and 30 ms [279, 280]. The difference reported after learning by the control, reward, and punishment groups in my study was approximately equal to the difference for the rewarded group (Wachter and colleagues: ~30 ms) and more than observed in the control and punishment groups in the study by Wachter and colleagues. Therefore, I am confident that participants were able to learn and express sequence-specific knowledge in all three feedback conditions.

4.4.1 Limitations

One aspect of the present study worth noting is the difficulty in comparing performance across tasks. This issue is discussed in Chapter 8.

One other limitation in the present study was the absence of the awareness test for the FTT. As discussed in Chapter 2, the SRTT has been previously used to investigate implicit sequence learning [157], while the FTT has not. In addition, because the trial structure of the SRTT consists of discrete units, this task is well suited to self-generation experiments. Considering the importance of explicit processes to other 'motor' tasks, such as visuomotor adaptation (e.g., [281]),

future research might consider the extent to which implicit and explicit processes are engaged during FTT performance.

In addition, regarding conscious recall, the awareness testing was conducted 30 days after training. It is possible that participants may have developed conscious sequence knowledge but had forgotten it by the time I measured awareness. The delay this was necessary to preserve the integrity of the delayed test-probes. To address this limitation, I investigated how conscious intent modulates the impact of feedback during SRTT learning in the study presented in Chapter 5.

4.5 Conclusions

Overall, I have demonstrated that punishment modulates on-line performance in a task-specific manner and shown that neither reward nor punishment modulates long-term retention of skill memories. These findings cast doubt on the commonly held hypothesis that reward is ubiquitously beneficial to memory and suggest that the interaction between feedback and learning should be better understood before feedback can be fully exploited in clinical situations. In the following Chapter, I further investigate the impact of valenced feedback on SRTT learning by investigating how the participant's intention to learn modulates the effect of reward and punishment.

5 Intentionality modulates the impact of feedback on performance

In the previous Chapter, I demonstrated that the impact of reward and punishment on performance during learning depends on task demands. The present study extends the work in Chapter 4 by investigating whether the impact of reward and punishment is consistent when participants are instructed to learn in a goal-directed fashion. In Experiment 1, participants performed the serial response time task with reward, punishment, or control feedback and were instructed to ignore the presence of the sequence, i.e., learn unintentionally. Experiment 2 followed the same design, but participants were instructed to intentionally learn the sequence. I found that punishment was beneficial to performance during learning only when participants learned unintentionally. Thus, in addition to task demands, the impact of feedback on performance is influenced by the goal of the learner.

5.1 Introduction

As described in Chapter 4, rewards and punishments play critical roles in shaping human behaviour during learning [76, 260, 282]. Yet despite the importance of valenced feedback for learning, the effects of reward and punishment on performance during motor skill learning have been largely unexplored [161]. In addition to the experiment presented in Chapter 4, only

Wachter and colleagues [152] have investigated this question in the context of motor sequencing using the SRTT. In both cases, punishment improved performance during learning [18, 152].

While these studies both show a convincing beneficial effect of punishment on performance, it is unclear whether this result generalizes to different learning contexts. In Chapter 4, I showed that task modulates the effect of punishment on performance. Even within a task, other factors may affect the nature of learning. For example, when engaging in learning in naturalistic settings, humans often have a conceptual understanding of their 'goal.' Basketball players, for example, will be aware that they must put the ball through the hoop, and they use this knowledge to guide their actions. This intention (putting in the ball in the hoop) shapes the learning process.

In the present Chapter, I use the terms 'implicit' and 'explicit' to refer to the quality of the information learned by the participant, and I use the terms 'unintentional' and 'intentional' to refer to the goal of the participant during performance. Despite its importance, intentionality has not been investigated in the context of skill learning and feedback. The SRTT is a good task for testing the influence of intentionality on learning. Using the SRTT, participants can be explicitly instructed to learn the sequence, or the knowledge of the sequence can be assessed post-hoc through the use of awareness tests [157, 270, 277]. In the experiment presented in Chapter 4, participants were not instructed about the presence of the sequence and largely demonstrated implicit acquisition when tested at 30-days post-learning (participants would have been excluded if they had demonstrated obvious explicit awareness of the sequence). In the study by Wachter and colleagues, participants were excluded if they achieved explicit

awareness of the presence of the sequence [152]. Thus, both studies left open the possibility that the effect of valenced feedback on performance may differ if participants acquire explicit knowledge or use an intentional strategy during learning.

Because of the central importance of intention to skill learning, in the present study I sought to investigate whether the intention to learn modulates the effect of valenced feedback on SRTT performance. In both experiments presented in this Chapter, participants performed the SRTT task augmented with reward, punishment, or control feedback. In the first experiment (the unintentional learning condition), I told participants that a sequence would be present but to ignore the sequence during training, and continue to press buttons as fast and accurately as possible. In the second experiment (the intentional learning condition), I told participants to try to learn the sequence. Based on the results presented in Chapter 4 demonstrating that both reward and punishment enhanced early learning compared to control feedback, I was particularly interested in the difference between the intentional and unintentional learners during the early learning period.

5.2 Methods

This study investigated the question of intentionality using two separate experiments. Experiment 1 investigated whether either reward or punishment benefitted SRTT performance when participants were instructed to learn unintentionally, see below. Experiment 2 determined whether this effect was still present when subjects were instructed to learn the sequence intentionally.

Specific details about the intentionality manipulation are explained below. The experimental timeline is presented in Figure 5.1a.

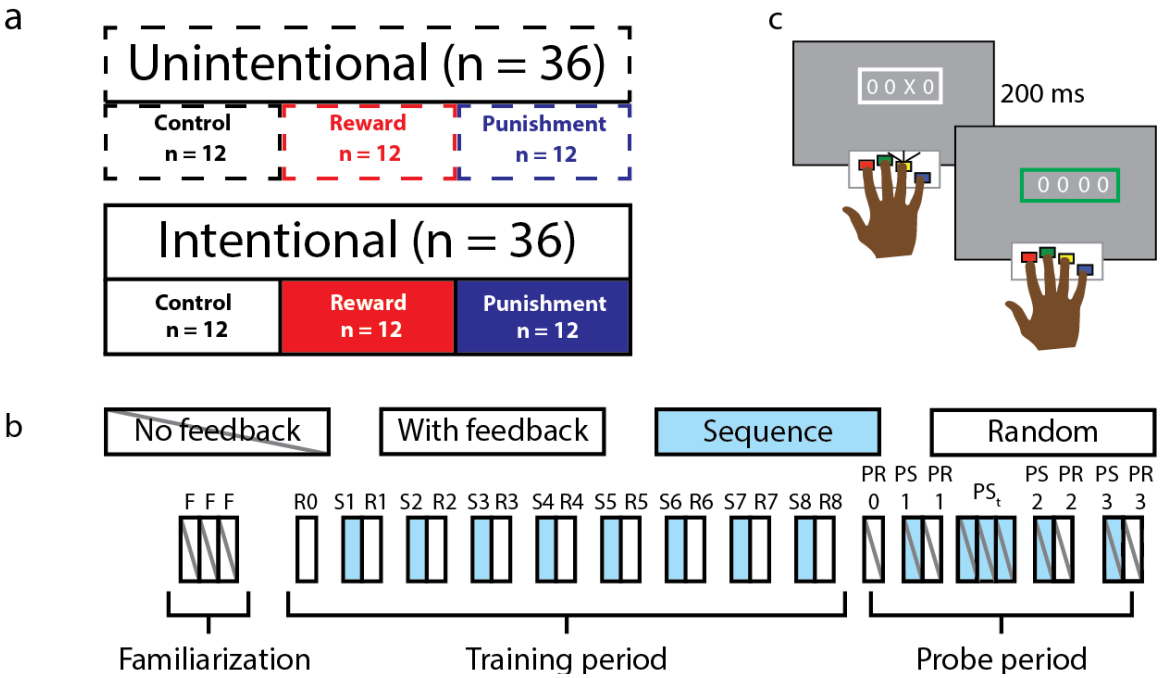


Figure 5.1. Experimental design. (a) The present study consisted of two experiments, which differed based on the instruction given to the participants. Participants performed the SRTT and were asked to either intentionally learn the sequence or to try to ignore the sequence. (b) The SRTT took place over 30 blocks. The first three blocks of the task (familiarization blocks [F]) contained no feedback and the stimulus appeared according to a random-sequence. The training period began at the fourth block (random-sequence [R]), at which point feedback was introduced. Beginning in the fifth block, the block order altered such that fixed- and random- sequence blocks were presented in an alternating order. Participant's RTs during in each fixed-sequence [S] block to the subsequent random-sequence block indicated skill knowledge expression. After the 17th block (eighth fixed-/-random doublet), the immediate retention probe began. During the immediate retention probe, participants performed a random-sequence block without feedback. They then performed a fixed-sequence block followed by a random block, and comparison between these two blocks indicated retention. Participants then performed four consecutive fixed-sequence blocks, followed by a random block. Finally, participants performed a fixed- and fixed-random block. c. In both experiments, to perform the SRTT with feedback. This was implemented in the same manner as Chapter 4.

The serial reaction time task (SRTT) was presented over 30 blocks. Each block featured 48 trials (Figure 5.1b). During some blocks ("fixed-sequence blocks" [S]) the stimulus would appear according to a repeating pattern. During other

periods, the appearance of the stimulus was randomly determined ("random-sequence blocks" [R]). During the training period, fixed- and random-sequence blocks were presented in an interleaved fashion [279]. I chose this design to make gaining explicit-knowledge more challenging for participants, and so that I could assess the development of sequence knowledge (compared to general skill) continuously across the learning period.

To familiarise participants with the task, and establish their baseline level of performance, the task began with three random-sequence blocks without feedback ("familiarization blocks"). The training period began with a random block, after which point the blocks alternated between random and sequence blocks (total number of blocks = 17; number of sequence blocks = 8). The training period concluded with a random block. The difference in performance between an S-block and the subsequent R-block (e.g., S1 – R1, S2 – R2, and so on) was used to index sequence knowledge during the training period [157]. Immediately following completion of the training period, participants were tested for sequence memory without feedback during the probe period. The probe period consisted of 9 blocks (Probe fixed-sequence [PS]; Probe random-sequence [PR]) without feedback in the following order: PR0 – PS1 – PR1 – PSt – PS_t – PS_t – PS2 – PR2 – PS3 – PR3. Sequence knowledge was indexed by comparing the initial sequence block (PS1) to the subsequent random block (PR1). Prior research using visuomotor adaptation tasks has suggested that subjects trained with punishment relearn at a faster rate [151]. I sought to determine whether this phenomenon occurred in sequence learning by including four fixed-sequence blocks after the first sequence knowledge test. The probe period concluded with a second

sequence knowledge test to determine whether the groups differed in the sequence knowledge after the prolonged sequence period.

To test the impact of reward and punishment on skill learning, participants were randomized into one of three Feedback groups: reward, punishment, or uninformative (control). During the feedback period, reward, punishment, or control feedback was provided based on the participant's ongoing performance.

5.2.1 Participants

72 participants (Experiment 1: 18 female, age = 24.4±3.1 [mean ± standard deviation]; Experiment 2: 17 female, age = 23.3±2.4) participated in this experiment. 36 subjects were assigned to each study. There was no difference between the average age of the participants in Study 1 and 2 ($t_{(70)} = 1.78$, $p = 0.0784$). All participants were right-handed, free from neurological disorders, and had normal or corrected-to-normal vision. All participants gave informed consent and the study was performed with approval from the Oxford Central University Research Ethics Committee (R44415/RE001). Participants were compensated using the local currency (£).

5.2.2 Serial reaction time task

The SRTT was implemented in a manner similar to Chapter 4. The methods are briefly represented here.

The version of the SRTT used here adds feedback to the traditional implementation (Figure 5.1c). At the beginning of each block participants were presented with four "O"s, arranged in a line at the centre of the screen. These stimuli were presented in white on a grey background. A trial began when one of the "O"s changed to an "X". Participants were instructed to respond as quickly and accurately as possible, using the corresponding button, on a four-button response

device held in their right hand. The "X" disappeared once the subject made a response, followed by a 200 ms fixed inter-trial interval, during which time the four "O"s were displayed.

A block consisted of 48 trials. During fixed-sequence blocks, the stimuli appeared according to a fixed 12-item sequence repeated 4 times (e.g., 3-4-1-2-3-1-4-3-2-4-2-1). Unlike Chapter 4, each fixed block began at the same position in the sequence. In the random blocks, the stimuli appeared according to a randomly generated sequence, without repeats on back-to-back trials, so, for example, participants would never see the triplet 1-1-2.

Between each block, participants saw the phrase "Nice job, take a breather." After five seconds, a black fixation-cross appeared on the screen for 25 seconds. Five seconds before the next block began, the cross turned blue to cue the participants that the block was about to start.

The first button press made after stimulus presentation was considered the participant's response. Only correct trials were considered for analysis of RTs. Trials with RTs less than 150 ms and greater than 800 ms were excluded from the analysis, which is consistent with prior work using the SRTT task.

5.2.3 Intentionality manipulation

In Experiment 1, after the familiarization period, participants were given the following instructions, "During certain periods of time, you may have the feeling that the stimulus is following a fixed sequence. Please ignore the sequence and continue to perform as fast and accurately as possible."

In Experiment 2, participants received the following instructions, "During certain periods of time, you may have the feeling that the stimulus is following a

fixed sequence. Please do your best to learn the sequence, while also performing as fast and accurately as possible."

Participants were not given any further instructions about the nature of the task and were not permitted to ask any additional questions about the sequence. Thus, I refer to the participants in Experiment 1 as "unintentional learners" and those in Experiment 2 as "intentional learners."

5.2.4 Feedback

The feedback procedure implemented in this study matched the procedure in Chapter 4. All participants were paid a base remuneration of £15 for participating in the study. At the start of the training period, participants were informed they could earn more money based on their performance. Participants were randomly assigned into the reward, punishment, or control feedback groups. The presence of reward or the absence of punishment was based on participant's performance. In both versions of the SRTT, an initial criterion RT was defined based on the participant's median performance during the final familiarization block. As participants progressed through training, this criterion was re-evaluated after each block, to encourage continuous improvement. In the reward group, the feedback indicated that the participant's performance was improving. In the punishment group, the feedback indicated they were getting worse. The control groups received feedback on 50% of trials. Participants in the control group were made aware that the feedback was not meaningful. I considered the reward and punishment control groups together in the analyses, as is typical in these studies [17, 152].

Performance was defined as the accuracy (correct or incorrect) and RT of a given trial. Feedback was given on a trial-by-trial basis and was indicated to the

participant by the white frame around the stimulus changing to green (reward) or red (punishment). In the reward group, the participants were given feedback if their response was accurate and their RT was faster than their criterion RT, which indicated that they earned money (£0.02 from a starting point of £0) on that trial. In the punishment group, participants were given feedback if they were incorrect, or their RT was slower than their criterion, which indicated that they lost money (£0.02 deducted from a starting point of £20) on that trial. Participants in the control group saw red or green colour changes, respectively, on 50% of trials that were randomly determined (and therefore unrelated to their performance). To ensure they were adequately motivated, participants in the control group were told that they would be paid based on their speed and accuracy. Importantly, to control for the motivational differences between gain and loss, participants were not told the precise value of a given trial. This allowed me to assess the hedonic value of the feedback, rather than the level on a perceived-value function. Between blocks, for the reward and punishment groups, the current earning total was displayed (e.g., "You have earned £5.00"). control participants saw the phrase, "You have earned money." The initial criterion RT was calculated as median performance in the first familiarization block. After each block, the median + standard deviation of performance was calculated and compared with the criterion. If this test criterion was faster than the previous criterion, the criterion was updated. Only correct responses were considered when establishing the criterion RT.

5.2.5 Conscious recall

To assess conscious recall, I used both verbal report and implemented the process dissociation procedure in the same manner described by [270]. Briefly,

after completion of the probe blocks, participants performed 3 sets of 100 button presses. First, I asked participants to generate 100 button presses randomly with no repeats. At this point, the presence of the sequence was revealed, and participants were then asked to input 100 button-presses while trying to repeat as much of the sequence as they could recall. Participants were specifically instructed that if they thought they knew any of the sequence, or several chunks of the sequence, that they should input those pieces as many times as they could. Participants were then asked to generate 100 button presses where they tried to exclude the sequence as much as possible. I added the 'random generation' test for the process dissociation procedure, based on my observation in Chapter 4 that there was a tendency for rewarded participants to perform faster on small portions of the fixed-sequence trials that appeared in the random-sequence blocks, leading to my hypothesis that reward may prime participants to perform the sequence more readily even during random button pressing. I also added the exclusion condition to assess conscious 'control' over excluding the sequence, as has been implemented previously in the literature [270, 283, 284]. The area under the curve for each number of sequential items (triplets, quadruplets, quintuplets through dodecs) included in the 100 button presses generated by the subject was compared to chance, which was estimated via 10,000 Monte-Carlo simulations of 100 non-repeating items with the probability of occurrence uniformly distributed across the 4 buttons. The area under the curve for each group was compared using a one-way ANOVA with Group (Reward/Control/Punishment) as a factor.

For the verbal report, I asked participants to report to me as much of the sequence as they knew. I added the verbal report because, unlike the data presented in Chapter 4, I only tested participants with an immediate probe, so any

explicit knowledge was more likely to be verbalized than after a prolonged delay. The total number of correct sequential items reported was compared using a one-way ANOVA with Group (Reward/Control/Punishment) as a factor.

5.2.6 Statistical analysis

Data analysis was conducted using JASP (Version 0.8.4, for Mac). To ensure there were no initial differences between the groups, the mean RT during the familiarization blocks were compared using a one-way ANOVA with Group (Reward/Punishment/Control) as a factor. To evaluate the impact of reward and punishment during the training period, RTs were compared across the groups using a repeated measures ANOVA with Group (Reward/Punishment/Control), Sequence-type (Fixed-/Random- sequence), and Block number (1-8) as factors. For the probe period, the first and second sequence knowledge tests were compared using a repeated measures ANOVA with Group (Reward/Punishment/Control), Sequence-type (Fixed-/Random- sequence), and Test (Test 1/Test 2) as factors. To test for differences in relearning, the four consecutive sequence blocks during the probe period were compared using a repeated measures ANOVA model with Group (Reward/Punishment/Control) and block (PS2-5) as factors.

To formally test whether the learning rate differed between the unintentional and intentional experiments, RTs were compared across the groups using a repeated measures ANOVA with Experiment (Intentional/Unintentional), Group (Reward/Punishment/Control), Sequence-type (Fixed-/Random-sequence), and Block number (1-8) as factors. Because I was specifically interested in learning rate, to control for total sequence knowledge I included the

performance during Probe 1 as a covariate. Results were qualitatively similar when the covariate was not considered.

Post-hoc tests were conducted using t-tests, with multiple comparisons correction using the Bonferroni-Holm method where appropriate.

5.3 Results

5.3.1 Experiment 1 – Unintentional learning

5.3.1.1 Punishment improves sequence knowledge expression during unintentional learning

The groups did not differ in their RT during the familiarization period (One-way ANOVA: $F_{(2,33)} = 1.372$, $p = 0.27$). I first sought to determine whether training with reward or punishment influenced the acquisition and expression of sequence knowledge (indexed by comparing RTs during fixed- and random-sequence blocks) during the training period. To this end, I compared the baselined RT (i.e., improvement compared to the three familiarization blocks to control for individual differences in RT) during the training period (S1 through R9) using a repeated measures ANOVA with Group (Reward/Punishment/Control), Sequence-type (Fixed-/Random- sequence), and Block (1-8) as factors (Figure 5.2a).

This analysis revealed that the difference in the RT between the fixed- and random-sequence blocks varied among the feedback groups over the course of learning (three-way interaction, Group x Sequence-type x Block: $F_{(7,231)} = 1.741$, $p = 0.049$). Follow-up analysis revealed the punishment group evidenced greater sequence knowledge during the first block compared to the reward and control groups (Punishment vs Reward: $t_{(23)} = 2.77$, $p_{corr} = 0.026$; Punishment vs Control:

$t_{(23)}$ = 2.79, p_{corr} = 0.026; Figure 5.2b) . The punishment group also evidenced greater sequence knowledge at the final block of training with feedback compared to the control group but not to the reward group (Punishment versus Control: $t_{(23)}$ = 2.81, p_{corr}= 0.022; Punishment versus Reward: $t_{(23)}$ = -1.013, p_{corr}= 0.19). There were no significant differences in performance between the feedback groups at any other block. There was no difference between the reward and control groups at the beginning ($t_{(23)}$ = 0.02, p_{corr}= 0.98) or end of training ($t_{(23)}$ = 1.0, p_{corr}= 0.57), suggesting that punishment, rather than reward, is beneficial to performance during unintentional sequence learning.

In addition to the difference in performance across the feedback groups that differed by block-type, the overall RT differed across the feedback groups over time (Group x Block interaction: $F_{(11,183)}$ = 2.47, p = 0.006). Post-hoc analysis revealed that the participants in the punishment group performed significantly faster than the control group during Block 3 ($t_{(23)}$ = 2.47, p_{corr} = 0.046) and Block 7 ($t_{(23)}$ = 2.67, p_{corr} = 0.033). The reward group was not significantly faster than the control group during any block (all ps > 0.05).

There was no effect of feedback group on accuracy during the period with feedback (Main effect of Group group: $F_{(2,33)}$ = 2.34, p = 0.112; Group x Block: $F_{(14,231)}$ = 0.97, p = 0.478; Group x Sequence-type: $F_{(2,231)}$ = 0.58, p = 0.56; Group x Sequence-type x Block: $F_{(14,231)}$ = 1.17, p = 0.30).

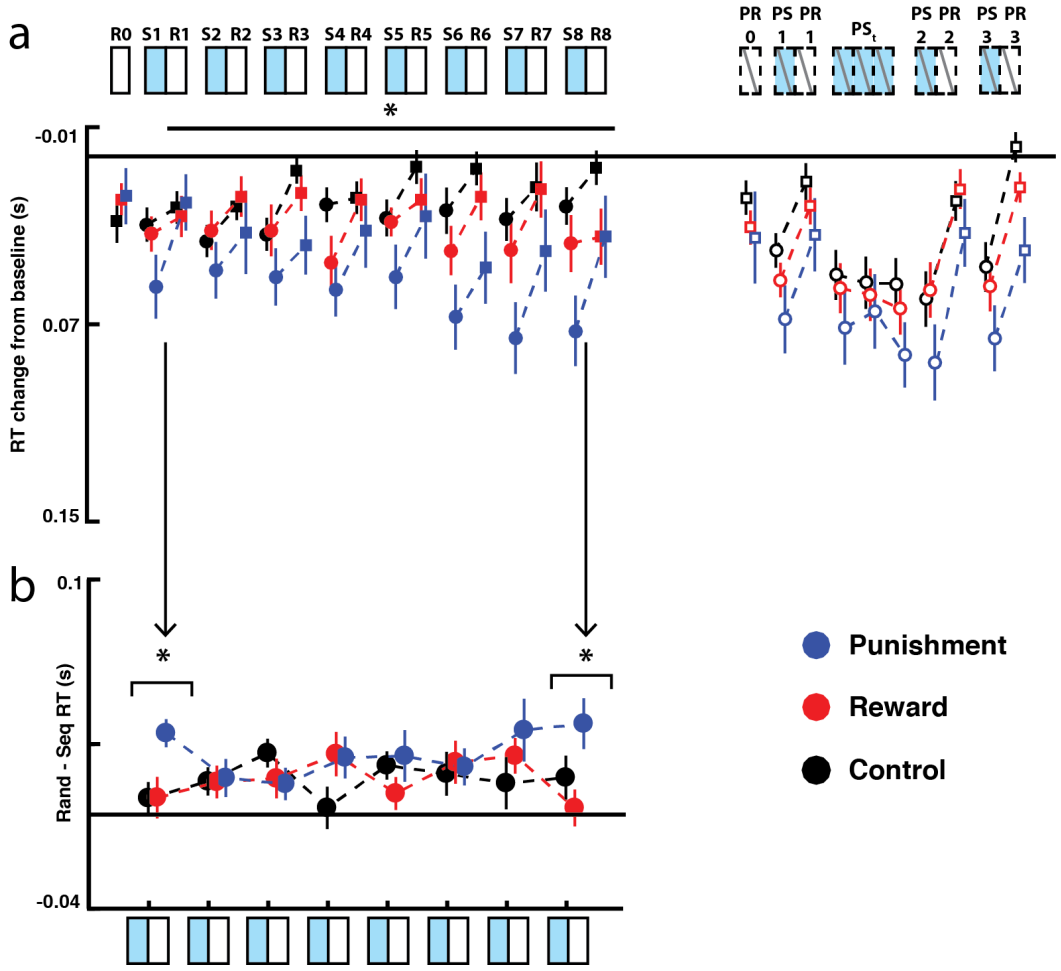


Figure 5.2. Punishment improves unintentional sequence learning. (a) The subject's performance during the training and immediate test conditions relative to their performance during the familiarization blocks. (b) Difference between fixed- and random-sequence RTs over the learning period. The punishment group evidenced more sequence knowledge during the first block of training with feedback, as well as during the final block of training with feedback. Asterisks indicate significant difference ($p < 0.05$). Error bars show SEM. R = random-sequence, S = fixed-sequence, PR = probe random-sequence, PS = probe fixed-sequence.

5.3.1.2 Short-term retention is not impacted by valenced feedback during unintentional learning

I next sought to determine whether valenced feedback impacted short-term sequence-knowledge expression. First, to examine immediate skill performance, I compared RTs during the first fixed- and random-sequence blocks presented immediately after training in the probe period. I compared the RTs during this probe using a repeated measures ANOVA with Group

(Reward/Control/Punishment) and Sequence-type (Fixed-/Random-sequence) as factors. All groups expressed skill knowledge during the immediate probe (Main effect of Sequence-type, $F_{(1,33)}$ = 39.01, p < 0.001; Fixed- vs Random-sequence: $t_{(35)}$ = 6.40, p < 0.001). There was no effect of Group on skill knowledge expression immediately after learning (Main effect of Group: $F_{(2,33)}$ = 1.10, p = 0.36; Group x Sequence-type: $F_{(2,33)}$ = 0.616, p = 0.55).

Punishment has previously been shown to increase relearning rate in visuomotor adaptation tasks, and I sought to determine whether this phenomenon also occurred in sequence learning. To examine this question, I compared the RTs during the four-consecutive fixed-sequence blocks in the probe period (retraining-probe) using a repeated measures ANOVA with Group (Reward/Punishment/Control) and Block (PS2-PS5) as factors. There was no effect of Group ($F_{(2,33)}$ = 1.02, p = 0.37), Block ($F_{(3,99)}$ = 2.098), p = 0.11), or Group x Block interaction ($F_{(6,99)}$ = 0.95, p = 0.46). There was also no difference between the groups when challenged with a random-sequence block after the retraining probe, assessed by comparing the RT in the final fixed-sequence block to the subsequent random-sequence block (Group x Sequence-type: $F_{(2,33)}$ = 0.577, p = 0.57). While all groups demonstrated learning (retraining probe – Main effect of Sequence-type: $F_{(1,33)}$ = 68.43, p < 0.001), there was no difference in skill knowledge across the groups in the final skill knowledge probe (Group x Sequence-type: $F_{(2,33)}$ = 0.577, p = 0.57). In the final retention probe, the punishment group was significantly faster than the control group (Main effect of Group ($F_{(2,33)}$ = 3.942, p < 0.029); Punishment versus Control: $t_{(23)}$ = 2.762, p_{corr} = 0.028, Punishment versus Reward: $t_{(23)}$ = 1.817, p_{corr} = 0.156, Reward versus Control: $t_{(23)}$ = 0.945, p_{corr} = 0.352). However, this was not specific to the trained

sequence, as all groups showed learning (Main effect Sequence-type $F_{(1,33)}$ = 57.512, $p < 0.001$; Sequence versus Random: $t_{(35)} = 7.702$, $p_{corr} < 0.001$), and there was no difference between the feedback groups in terms of sequence knowledge expression (Sequence-type x Group $F_{(2,33)} = 0.468$, $p = 0.63$).

5.3.1.3 Feedback does not alter awareness during unintentional learning

Lastly, I sought to determine whether the addition of reward or punishment altered conscious awareness during unintentional skill learning. To assess the degree of implicit/explicit knowledge, I utilized two different measures. First, to assess any purely implicit bias, before the presence of the sequence was declared, I asked participants to perform a series of button presses and try to be as random as possible. I then had participants performed the process dissociation procedure, wherein they were asked to include or exclude the sequence during a series of button presses using the method described by [270]. One subject was removed from the analysis due to data collection failure. I calculated the area under the curve which describing the number of sequential items generated (doubles, triples, quads, etc) for each set of button pressing and compared across the groups on this measure using a one-way ANOVA (Figure 5.3).

Using this procedure, I found no difference in sequence inclusion on any set of presses between the feedback groups (Random: $F_{(2,32)} = 0.528$, $p = 0.95$; Inclusion: $F_{(2,32)} = 0.400$, $p = 0.68$; Exclusion: $F_{(2,32)} = 0.233$, $p = 0.79$).

As a second approach, participants indicated knowledge via verbal report. For this test, I asked participants to verbally recall as much of the sequence as they were confident in producing. The number of sequential triplets included was

compared across the groups using a one-way ANOVA. The feedback groups did not differ in the number of sequential items reported ($F_{(2,33)} = 1.62$, $p = 0.21$).

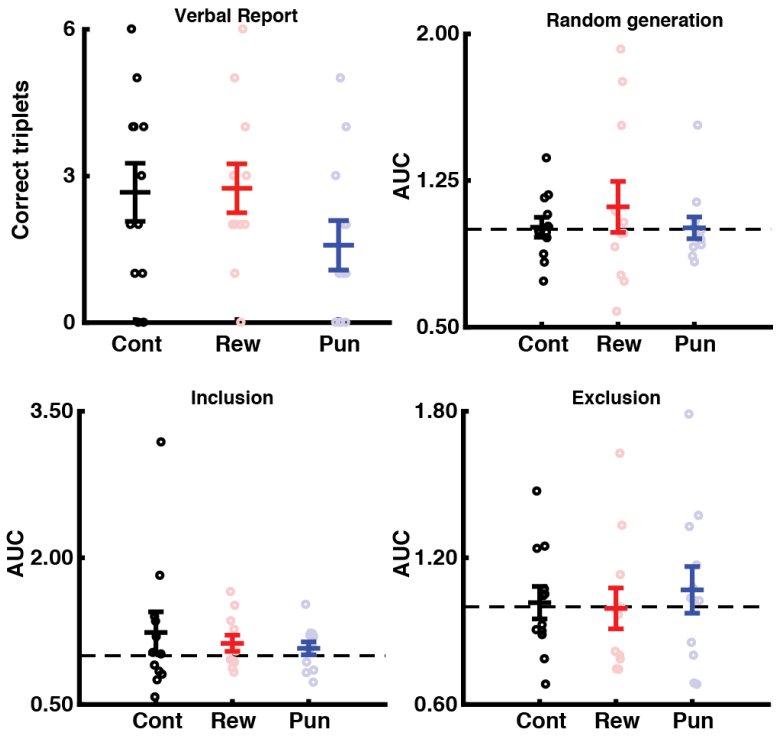


Figure 5.3. Feedback did not impact awareness when learning unintentionally. There was no difference between the Feedback groups with respect to explicit knowledge formed during the unintentional learning condition based on Verbal report, as well as process dissociation procedure inclusion, exclusion, or free generation rates. Error bars show SEM. Colors: Black = control, red = reward, blue = punishment

5.3.1.4 Unintentional learning summary

In summary, when learning unintentionally, I found that participants training with punishment demonstrated significantly more skill knowledge during the acquisition period than participants training with reward or control feedback. The feedback given during the training period did not influence their ability to express sequence knowledge immediately after training, once the feedback was removed, and did not systematically modulate the implicitness of the sequence acquisition.

5.3.2 Experiment 2 – Intentional learning

5.3.2.1 Punishment does not benefit intentional sequence learning

To determine whether punishment similarly benefitted intentional sequence learning, I performed the same experiment described above, but instructed participants to intentionally learn the sequence. Data from this experiment are shown in Figure 5.4a. There was no difference between the feedback groups during the training period ($F_{(2,33)}$ = 1.094, p = 0.34). When instructed to learn intentionally, all groups showed sequence knowledge improvement over time (Sequence-type x Block interaction: $F_{(7,231)}$ = 4.81, p < 0.001; Block 1 versus Block 8, $t_{(35)}$ = 3.781, p < 0.001), but the groups did not differ in overall performance or sequence knowledge (Main effect of Group: $F_{(2,33)}$ = 1.13, p = 0.34; Group x Block: $F_{(14,231)}$ = 0.903, p = 0.56; Group x Sequence-type: $F_{(2,33)}$ = 0.244, p = 0.79; Group x Sequence-type x Block: $F_{(7,231)}$ = 1.36, p = 0.18; Figure 5.4). Thus, in contrast to the benefit of punishment found in unintentional learning, when instructed to learn intentionally, punishment did not express sequence knowledge earlier the reward and control groups.

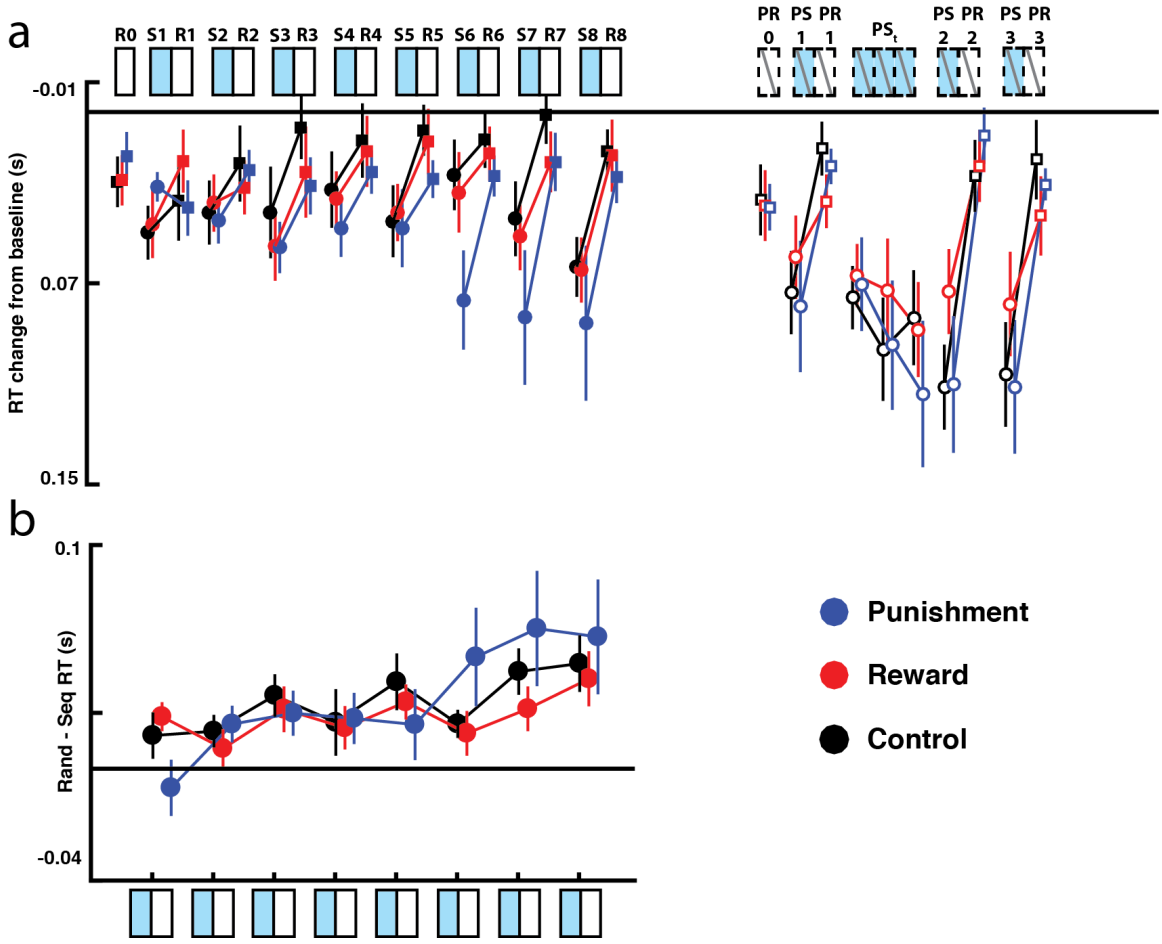


Figure 5.4. Feedback does not impact performance during intentional sequence learning. (a) All groups showed robust learning when instructed to learn the sequence (b) Difference in RTs during fixed- and random-sequence blocks across the learning period. All groups showed equivalent development of sequence knowledge. Error bars show SEM. R = random-sequence, S = fixed-sequence, PR = probe random-sequence, PS = probe fixed-sequence.

5.3.2.2 Feedback does not alter retention when learning intentionally

I next investigated whether feedback modulated retention or relearning after intentional sequence learning. All groups expressed skill knowledge during the immediate probe (Main effect of Sequence-type, $F_{(1,33)} = 26.00$, $p < 0.001$; Fixed- vs Random-sequence: $t_{(35)} = 5.09$, $p < 0.001$). However, there was no effect of Group on skill knowledge expression immediately after learning (Main effect of Group: $F_{(2,33)} = 1.09$, $p = 0.35$; Group x Sequence-type: $F_{(2,33)} = 0.26$, $p = 0.80$).

When I examined the rate of relearning, I found that all groups showed improvement during the relearning period (Main effect of Block: $F_{(3,99)} = 9.105$, $p < 0.001$). However, there was no effect of feedback on relearning (Main effect of Group: $F_{(2,33)} = 0.68$, $p = 0.51$; Group x Block interaction: $F_{(6,99)} = 0.95$, $p = 0.46$). Similar to the unintentional learning, all groups showed sequence knowledge after relearning (Main effect of Sequence-type: $F_{(1,33)} = 68.475$, $p < 0.001$), but there was no impact of Group on the sequence knowledge expressed immediately after relearning (Group x Sequence-type: $F_{(2,33)} = 0.50$, $p = 0.61$). In the final retention probe, all groups showed learning (Main effect Sequence-type $F_{(1,33)} = 30.97$, $p < 0.001$), but there was no effect of feedback on performance or evidenced knowledge (Main effect of Group: $F_{(2,33)} = 0.79$, $p = 0.47$; Sequence-type x Group $F_{(2,33)} = 1.05$, $p = 0.36$).

5.3.2.3 Awareness was not modulated by feedback when learning intentionally

Finally, I sought to determine whether the degree of explicit awareness after intentional sequence learning was modulated by feedback (Figure 5.5). Using the process dissociation procedure described above, I found no difference in sequence inclusion on any set of presses (Random: $F_{(2,32)} = 0.35$, $p = 0.70$; Inclusion: $F_{(2,32)} = 0.1.48$, $p = 0.24$; Exclusion: $F_{(2,32)} = 0.7$, $p = 0.50$). The feedback groups also did not differ in the number of sequential items verbally reported ($F_{(2,33)} = 0.36$, $p = 0.70$).

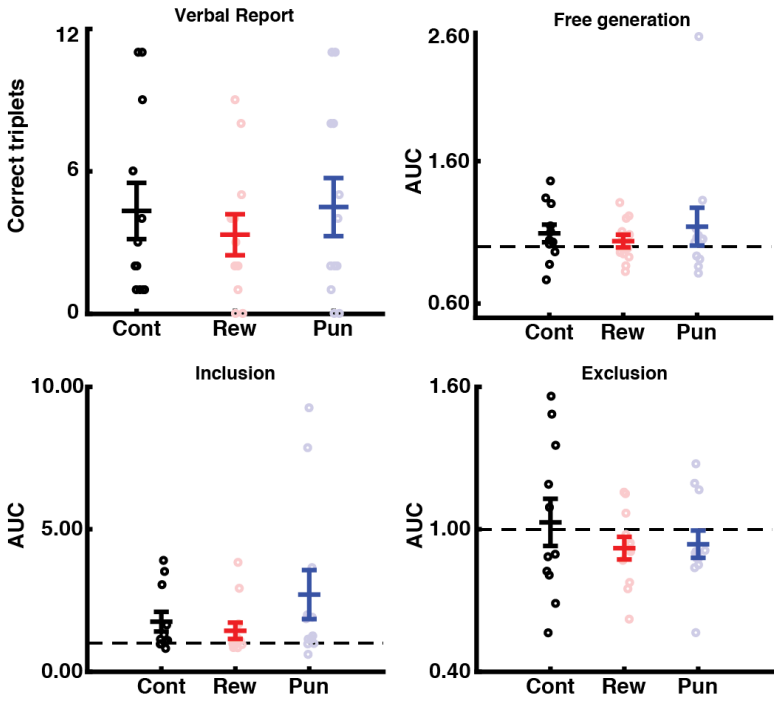


Figure 5.5. There was no difference between the feedback groups with respect to explicit knowledge formed during intentional sequence learning. Error bars show SEM. Colors: Black = control, red = reward, blue = punishment.

5.3.3 Punishment specifically benefits unintentional learning

The present study sought to investigate the influence of intentionality on learning with valenced feedback. Therefore, I directly compared the performance during learning from Experiments 1 and 2 using a repeated measures ANOVA with Intentionality (Intentional/Unintentional), Group (Reward/Punishment/Control), Sequence-type (Fixed-/Random- sequence), and Block (1-8) as factors. To isolate learning-rate from overall learning, performance during the first test probe was included as a covariate.

This analysis confirmed that intentionality did influence the effect of valenced feedback on sequence learning rate (4-way interaction, Intentionality x Group x Sequence-type x Block; $F_{(14,455)} = 1.835$, $p = 0.032$). Results were qualitatively similar when not controlling for performance during Probe 1 ($F_{(14,462)}$

= 1.657, p = 0.06). Follow-up tests revealed that intentionality specifically modulated the influence of punishment on early learning (Figure 5.6). The punishment group learning unintentionally evidenced more sequence knowledge during Block 1 compared to punishment group learning intentionally (Unintentional Punishment vs Intentional Punishment: $t_{(22)}$ = -3.241, p = 0.0038). The unintentional punishment group did not perform significantly faster than the intentional reward or intentional control groups (Unintentional Punishment vs Intentional Reward: $t_{(22)}$ = 1.350, p = 0.19; Unintentional Punishment vs Intentional Control: $t_{(22)}$ = 1.733, p = 0.10). Intentionality did not influence learning during the Block 1 in either the reward or the control groups (Unintentional Control vs Intentional Control: $t_{(22)}$ = 0.658, p = 0.51; Unintentional Control vs Intentional Reward: $t_{(22)}$ = 0.583, p = 0.56; Unintentional Reward vs Intentional Reward: $t_{(22)}$ =1.531, p = 0.14; Unintentional Reward vs Intentional Control: $t_{(22)}$ = 0.58, p = 0.566). The feedback groups did not differ at any other blocks. Thus, these data support the conclusion that intentionality specifically modulates the impact of punishment on early learning.

With respect to awareness measures, intentional learning fostered significantly greater explicit awareness, as evidenced by a greater number of sequence items generated during the verbal report ($F_{(1,66)}$= 6.452, p = 0.013; Unintentional v Intentional: $t_{(68)}$= 2.54, p = 0.013) and inclusion rate during the process dissociation procedure ($F_{(1,66)}$= 6.39, p = 0.014; Unintentional v Intentional: $t_{(68)}$= 2.50, p = 0.14). There was no main effect or interaction with feedback group in the awareness tests for either experiment (all ps > 0.3).

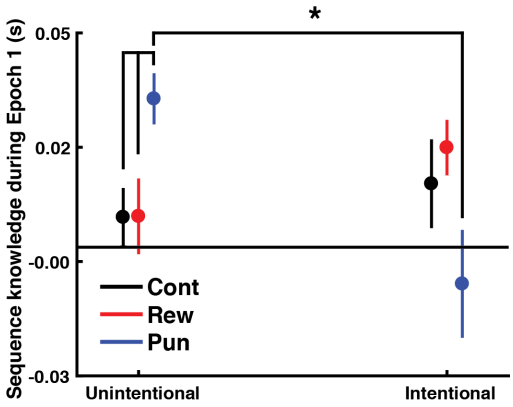


Figure 5.6. Comparison of early learning across the Feedback groups in both experiments. Punishment significantly benefitted early learning when instructed to learn unintentionally but not intentionally. Asterisks indicate a significant difference ($p < 0.05$). Error bars show SEM.

5.4 Discussion

In this study I investigated whether intentionality interacts with the effect of valenced feedback on performance during sequence learning. I found that punishment enhanced early sequence knowledge expression when participants learned unintentionally (Experiment 1). However, when learning intentionally, all groups performed equally well (Experiment 2). Direct comparison of the results of the two experiments revealed that punishment was beneficial specifically during the early phase of unintentional sequence learning. These results suggest that intentionality influences the impact of valenced feedback given during performance.

The effect of reward and punishment on performance during learning has previously been characterized with a simple heuristic: punishment being beneficial to performance, while reward benefits memory retention [17, 151, 152, 161]. In Chapter 4, I showed that this simple heuristic is not valid in all cases. Specifically, I demonstrated that the impact of valenced feedback depends on the task being performed: while punishment benefited performance on the SRTT

when learning unintentionally, punishment was detrimental to performance on a different motor learning task (the FTT). This experiment shows that, even when participants are performing the same task, altering the goal of the participant can influence the effect of reward and punishment, which further complicates the use of the simple heuristic even within the same task. Together, these two results suggest that valenced feedback interacts with motor skill learning in a complex manner.

The critical question raised by this study is why punishment benefits unintentional, but not intentional, learning. A critical component of the intentional learning is the added cognitive load introduced by intentionally learning the sequence. The 'cognitive load' may be due to several factors, but two examples are i) the added 'dual-task' of learning to maximizing the amount earned, and ii) the added performance monitoring during negative feedback that may cause 'choking under pressure' in procedural learning tasks.

It is possible that the goal of learning the sequence acts as a distractor task that competes with the primary task of pressing buttons quickly and accurately. While participants in the intentional and unintentional learning conditions share the goal of maximizing the total money earned, participants learning intentionally are also compelled to detect and learn a sequence, while participants learning unintentionally are not. The added cognitive load in the intentional condition, then, may be similar to dual-task learning. In dual-task learning, participants are forced to learn while simultaneously performing another task that competes for attention, for example counting tones [285-289]. Under dual-task conditions, sequence knowledge may take longer to be expressed [285-287], and, consistent

with the data from the present study, the detrimental effect of the added cognitive load may be more extreme when participants engage explicit strategies [289].

Although it should be noted that visuomotor adaptation is generally considered to be a separate process from motor-sequence learning, complementary evidence comes from the visuomotor adaptation literature. For example, it has been found that the beneficial effects of reward-feedback to retention in visuomotor adaptation tasks appear to depend on explicit processes, and a distractor-task that disrupts explicit processes completely abolishes learning when participants must use reward-based feedback to adapt [290, 291]. While studies of visuomotor adaptation have not explored the role of explicit processes when learning with punishment, this could be an interesting avenue for future research.

A second possible source of cognitive load that may account for the difference between intentional and unintentional learning is that the engagement of the intentional learning strategy in the context of negative feedback invokes meta-cognitive processes (e.g., performance monitoring) that interfere with learning [292-295]. One study has examined the role of performance monitoring in the context of the unintentionally learned SRTT where a control condition (no monitoring) was compared i) to an outcome-pressure condition, where participants were told that performance needed to improve to earn additional compensation, and ii) a monitoring-pressure condition, where participants were told that their performance was being videotaped and would be watched by researchers [292]. These researchers reported that while the outcome-pressure and control conditions did not differ in their performance during learning, the monitoring-pressure condition produced a significant detriment in performance.

Thus, if participants engage in greater performance monitoring during the intentional learning condition, they might exhibit a similar deficit in performance.

Within the unintentional learning condition, the effect of feedback differed over the course of the learning period. The punishment group showed greater sequence knowledge expression compared to the reward and control groups in the first and last learning blocks, but not during the middle blocks of the learning period. The time course of sequence knowledge expression may be due to the interleaving of fixed- and random-sequence blocks during the learning period. In the interleaved design, participants may undergo forgetting or unlearning of the sequence during the random-sequence blocks [296-298]. In this study, training with punishment might have enhanced early learning during the first sequence block, but also quickened adaptive forgetting during the subsequent random block, as instances of negative reinforcement may serve as an implicit signal to shift strategies or explore new behaviours (in this case, to forget any previously learned sequence knowledge). In contrast, participants training with reward may be less sensitive to slight changes in reinforcement frequency and be less likely to shift rules or undergo forgetting, leading to the development of sequence knowledge over a longer period of exposure. This pattern of behaviour is consistent with a win-stay, lose shift strategy [299, 300]. Notably, it has been shown that reward is not effective in shaping motor behaviours unless participants are aware of the manipulation being rewarded or if reward is too abundant [263, 301], probably due to participants not engaging in exploratory behaviours. Future work may also consider the psychological implications of reward and punishment given during learning, and how this interacts with behavioural variability and cognitive strategies.

The design implemented here precludes modelling the cognitive strategies, such as the win-stay lose-switch versus reinforcement learning strategies [300]. However, based on the role of dopamine in perseveration, i.e., the use of reinforcement-learning strategy, I speculate that punishment may promote participants to update their sequence-item knowledge more rapidly (effectively facilitating a win-stay lose-switch strategy), while reward may facilitate greater use of a reinforcement-learning strategy and integrate their sequence-knowledge over a longer period of time [302].

Besides cognitive load, one potential explanation for the lack of benefit of punishment to performance is that learning with punishment and explicit sequence learning utilize shared neural substrates. Explicit sequence learning is known to recruit the hippocampus [252], and the experiment I present in Chapter 6 provides evidence that training with punishment also engages MTL. It is possible that these two processes interact competitively specifically at the level of MTL, thereby reducing the expression of early sequence knowledge. However, this possibility may be interesting to explore in the future.

Reward did not benefit performance on the immediate retention probes in the intentional or the unintentional learning conditions, which effectively replicates the similar null result found in the experiments included in Chapter 4. Given the widely observed benefit of reward to visuomotor adaptation memory retention, it will be important to understand whether any specific conditions reliably foster reward-related memory benefits in the context of motor skill learning [151, 161, 266, 267].

One aspect of the present study worth considering is the extent to which the intentional learning condition should be considered a separate task to

unintentional learning. Implicit and explicit sequence learning are often considered together in the same studies and treated as equivalent conditions [150, 157, 277, 303], and so there is precedent in the literature to consider these types of manipulations as equitable variations of the same task. This assumption is primarily based on task demands; the two tasks are functionally equivalent, in the sense that the equivalent stimulus elicits the same response in both cases, and, importantly, the contingency to receive the reward or punishment is equivalent between the two tasks. Given these considerations, it is my opinion that, although these variations of the task undoubtedly engage two different cognitive processes, the intentional and unintentional learning conditions can be compared fairly as equivalent tasks.

5.4.1 Limitations

The design of the present study had several limitations. Firstly, this study lacked a delayed memory test. I did not use a delayed memory test because I was specifically interested in the early sequence knowledge advantage for punishment. The study by Wachter and colleagues [152] also only assessed immediate knowledge, and one goal here was to try to replicate their finding of the benefit of reward to immediate sequence knowledge expression. Secondly, because the blocks only featured 48 trials per block, the estimation of the median RT may have been less stable compared to the study in Chapter 4. I used this approach because I wanted to assess sequence learning more frequently. Finally, because the test probes featured a period of continuous fixed-sequence blocks, participants may have been able to gain more explicit awareness during the probe period. This makes interpreting the awareness tests less straight forward, as it is not clear when explicit knowledge would have been acquired.

5.5 Conclusions

In this Chapter, I replicated and extended the work presented in Chapter 4, showing that punishment enhances early sequence knowledge expression (on the SRTT) specifically when participants learn unintentionally. As is the case with visuomotor adaptation, the effect of valenced feedback in sequence learning may also depend on explicit processes. Combined with the work from Chapter 4, these results suggest that the effect of valenced feedback is nuanced and depends on the task being performed as well as the intentions and expectations of the learner. In the next Chapter, I investigate the impact of feedback on brain activity after training.

6 Differential impact of reward and punishment on functional connectivity after skill learning

In Chapters 4 and 5, I explored the influence that exogenous rewards and punishments exert on behavioural performance. I found that, while all participants showed equal memory for the learned skill, performance during learning differed dramatically depending on the demands of the learning environment, suggesting that different neural systems may be engaged depending on the feedback being given. The present chapter unpacks this question by examining the impact of feedback on resting-state functional connectivity change after training on the SRTT and the FTT. In both tasks, reward and punishment differentially affected functional connectivity change after training, but the specific regions depended on the task being learned. The behavioural data from this experiment was presented in Chapter 4.

6.1 Introduction

In Chapters 4 and 5, I presented evidence suggesting that reward and punishment differentially affect performance during learning, and that the nature of the effect (i.e., benefit or detriment) depends on the demands required by the task. While these studies provided a fuller picture of the effect of feedback on

performance, these behavioural studies did not address the neural systems that might support learning with feedback. One hypothesis is that the neural systems for feedback processing are common across tasks. This hypothesis specifically suggests that 'fast' and 'slow' learning systems underlie learning with punishment and reward [248]: punishment leads to the recruitment of fast, attention-dependent learning systems [e.g., prefrontal cortex, anterior insula, and MTL], while reward recruits slow learning systems (e.g., caudate or putamen via dopaminergic signalling) [152, 304]. In support of this hypothesis, functional imaging studies where fMRI is acquired concurrent with task performance have reported that reward increases caudate activity in a behaviourally-relevant manner [152, 304]. In contrast, punishment increases activity in the anterior insula [152, 305] and MTL [306, 307].

The experiments considered in Chapters 4 and 5 suggest that, rather than engage a common feedback processing network, reward and punishment may engage different brain regions depending on the demands of the task being performed. Studies examining the impact of reward and punishment generally use tasks that are statistical in nature, such as decision making [262], probabilistic learning [251, 308, 309], or SRTT learning [152]. Therefore, how reward and punishment affect brain activity when performing tasks that require different demands is not known.

To address this question, I examined the impact of valenced feedback on neural activity induced by SRTT training and FTT training collected during the experiment presented in Chapter 4. Before and after participants trained on the tasks, I collected 20-minutes of resting-state fMRI data to look for changes in functional connectivity induced by training (Figure 6.1a). Based on the study by

Wachter and colleagues, I hypothesized that when learning the SRTT, reward and punishment would affect neural activity in the slow and fast learning systems [152, 248]. I then tested whether any of the regions implicated in feedback processing during SRTT training generalized to the FTT.

When examining functional connectivity change after learning, I focused on PMC as the key region for evaluating post-encoding connectivity in both tasks. I based this choice on the well-documented role the PMC plays as a memory-encoding region for sequence and sensorimotor learning [15, 55-57]. In addition, PMC shows reward-related activity after movement [126].

6.2 Methods

In two experiments, participants were trained on either the SRTT or the FTT with reward, punishment, or uninformative feedback (Figure 6.1b-d). No participant was trained in both tasks. Before and after the training session, 20-minutes of resting-state fMRI was collected.

The present study consists of the resting-state fMRI data collected as part of the project presented in Chapter 4, and a detailed description of the behavioural protocol can be found in Section 4.2. In this Section, I provide details of the resting-state fMRI methodology and analysis.

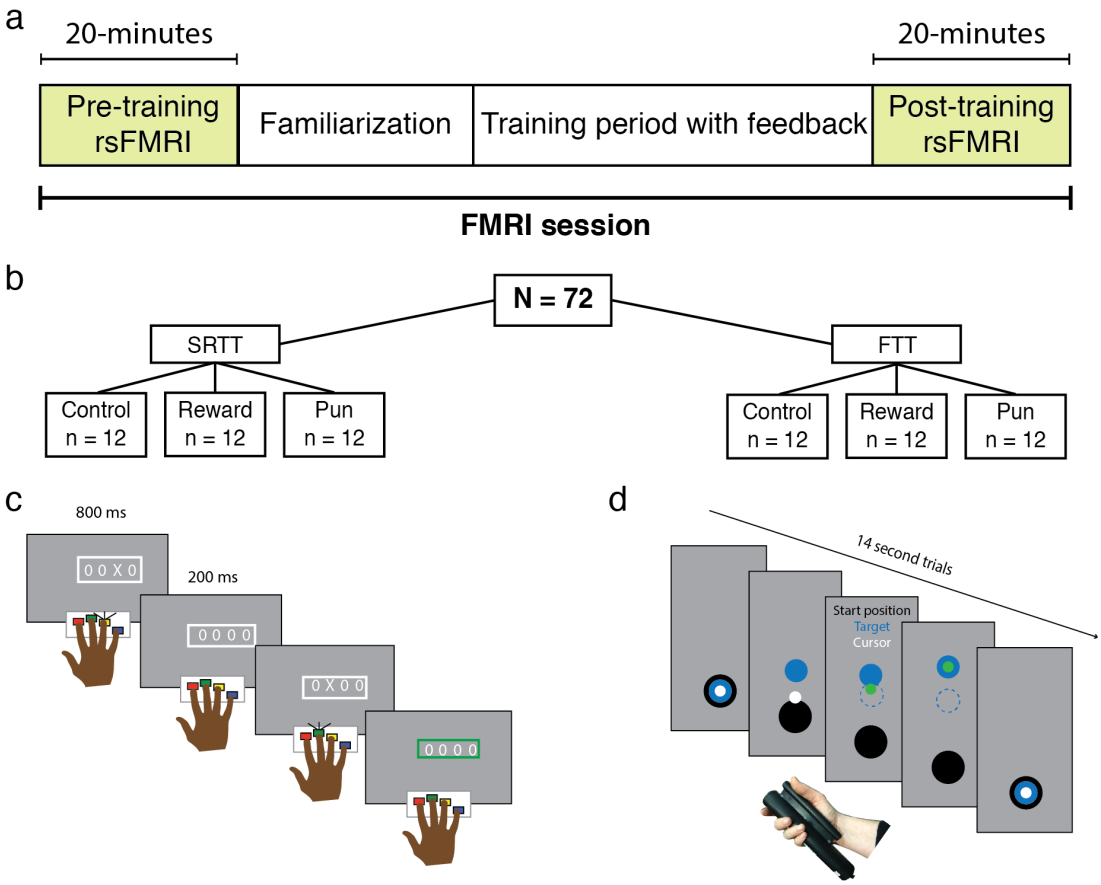


Figure 6.1. Experimental design and skill memory retention. (a,b). Participants underwent 20 minutes of resting-state fMRI before and after training on either the serial reaction time task (SRTT) or the force tracking task (FTT) while receiving reward, punishment, or control feedback. In the SRTT (c), participants responded to a cue appearing in one of four locations on a screen. In the FTT (d), participants modulated their grip force to track a moving target. In both tasks, the stimulus could follow either a random- or fixed-sequence, and skill memory was assessed by comparing performance during random- and fixed-trials.

6.2.1 Participants

78 participants (47 female, mean age = 25 years ± standard deviation 4.25) were recruited and participated in the study. All participants were right-handed, free from neurological disorders, and had normal or corrected-to-normal vision. All participants gave informed consent and the study was performed with NIH Institutional Review Board approval in accordance with the Declaration of Helsinki (93-M-0170, NCT00001360). Data from six individuals were removed

from the study due to inattention during training (defined as non-responsive or inaccurate performance on greater than 50% of trials; n=3) or inability to complete the imaging session due to discomfort or fatigue (n=3). This left 72 participants with complete data sets included in the analyses.

6.2.2 MRI acquisition

This experiment was performed on a 3.0T GE 750 MRI scanner using a 32-channel head coil (GE Medical Systems, Milwaukee, WI).

6.2.2.1 Structural scan

For registration purposes, a T1-weighted anatomical image was acquired (magnetization-prepared rapid gradient echo (MPRAGE), TR = 7 ms, TE = 3.4 ms, flip-angle = 7 degrees, bandwidth = 25.000 kHz, FOV = 24x24 cm^2, acquisition matrix = 256 x 256, resolution = 1 x 1 x 1 mm, 198 slices per volume). Grey matter, white matter, and CSF maps for each participant were generated using Freesurfer [236].

6.2.2.2 Echo-planar images

Multi-echo EPI scans were collected with the following parameters: TE = 14.9, 28.4, 41.9 ms, TR = 2, ASSET acceleration factor = 2, flip-angle = 65 degrees, bandwidth = 250.000 kHz, FOV = 24 x 24 cm, acquisition matrix = 64 x 64, resolution = 3.4 x 3.4 x 3.4 mm, slice gap = 0.3 mm, 34 slices per volume covering the whole brain. Respiratory and cardiac traces were recorded. Each resting-state scan lasted 21 minutes. The first 30 volumes of each resting-state scan were discarded to control for the difference in arousal that occurs at the beginning of resting-state scans. This left the final 20 minutes of rest in each scan for our

analysis. This procedure has been used in other studies where long-duration resting-state runs were collected [237].

6.2.3 Resting state fMRI preprocessing

Data were preprocessed using AFNI [198]. The time series for each TE was processed independently prior to optimal combination. Slice-time correction was applied (3dTShift) and signal outliers were attenuated (3dDespike)[192]. Motion correction parameters were estimated relative to the first volume of the middle TE (28.4 msec), and registered to the structural scan (3dSkullStrip, 3dAllineate). These registration parameters were then applied in one step (3dAllineate) and the data were resampled to 3 mm isotropic resolution. The three echoes were then combined using the procedure described in Chapter 2.

After optimal combination, I applied the basic ANATICOR [193] procedure to yield nuisance time series for the ventricles and local estimates of the BOLD-signal in white matter. All nuisance time-series (six parameters of motion, local white matter, ventricle signal, and 6 physiological noise regressors (RetroTS)) were detrended with fourth order polynomials. These time series, along with a series of sine and cosine functions to remove all frequencies outside the range (0.01-0.25 Hz) were regressed out in a single regression step (3dTproject). Time-points with motion greater than 0.3 mm were removed from the data and replaced with values obtained via linear interpolation in time. Data were transformed into standard space (@auto_tlrc) and smoothed with a 6mm FWHM Gaussian kernel. For group data analysis, a group-level grey matter by mask was created by calculating voxels in standard space determined to be grey matter in 80% of participants fMRI data analysis [175].

6.2.4 Premotor cortex functional connectivity analysis

I focused my analysis on the left PMC based on its well-described central role in sequence production in response to visual cues [310], sensorimotor learning, and sequence learning [55, 311]. Given that the participants were performing the task with their right hand, I further focused on the left PMC.

I defined left dorsal and ventral PMC based on a publicly available diffusion-MRI based parcellation of PMC [312]. The mean time series from both dorsal and ventral premotor cortices were extracted separately from each participant and each rest period, and the whole brain correlation maps (Pearson's r) for both dorsal PMC and ventral PMC during the pre- and post-training resting-state MRI scans were then calculated based on these time series. For each resting-state scan, GCOR [238] and first difference of motion (@1dDiffMag) were calculated and included as nuisance covariates for group analysis.

The resulting maps were then submitted to a LME model (3dLME) with ROI (PMd/PMv), Rest (pre-/post-), Group (Control/Reward/Punishment). The precise model fit at the group level in R-syntax (nlme) was 'Group*Rest*ROI+motion+gcor'. Group analysis maps were cluster-corrected for multiple comparisons to achieve a a= 0.05 using the ACF model in 3dClustSim ($p < 0.005$, $k = 54$).

To determine the consistency of the impact of feedback across the two tasks, the overlap of the significant clusters from the Rest x Condition interaction for the SRTT and FTT was calculated at both a conservative ($p < 0.005$, $k = 54$) and liberal ($p < 0.01$, $k = 54$) threshold.

6.3 Results

The full results of the LME model for both the SRTT and FTT are shown in Table 6.1. In the text I focus specifically on the effect of feedback on functional connectivity after training (Rest x Group interaction).

6.3.1 Reward and punishment evoked dissociable changes in PMC connectivity after SRTT training

In Experiment 1, following training on the SRTT, several regions showed a modulation of PMC functional connectivity that differed depending by feedback group (Rest x Group interaction; Figure 6.2). These included bilateral thalamus and striatum, right cerebellar vermis, SMA, bilateral MTL, and left inferior frontal gyrus. In order to examine the nature of the interaction, I extracted the estimated mean PMC functional connectivity change due to training from each cluster (Figure 6.2). This revealed a pattern that was distinct across the feedback groups. After training, the reward group had increased functional between PMC and thalamus and striatum, cerebellar vermis, and SMA, but PMC functional connectivity decreased with MTL and inferior frontal gyrus. The punishment group showed the opposite pattern: PMC connectivity with MTL and inferior frontal gyrus increased after training with punishment. The control group showed in intermediate pattern between the reward and punishment groups. In the control group, PMC connectivity with the thalamus, striatum, and cerebellar vermis increased after training, but PMC connectivity with SMA, MTL, and inferior frontal gyrus decreased after training. Thus, in the context of the SRTT, there is evidence multiple sets of brain regions change functional connectivity with PMC

after training, and these can be dissociated by augmenting training with reward or punishment.

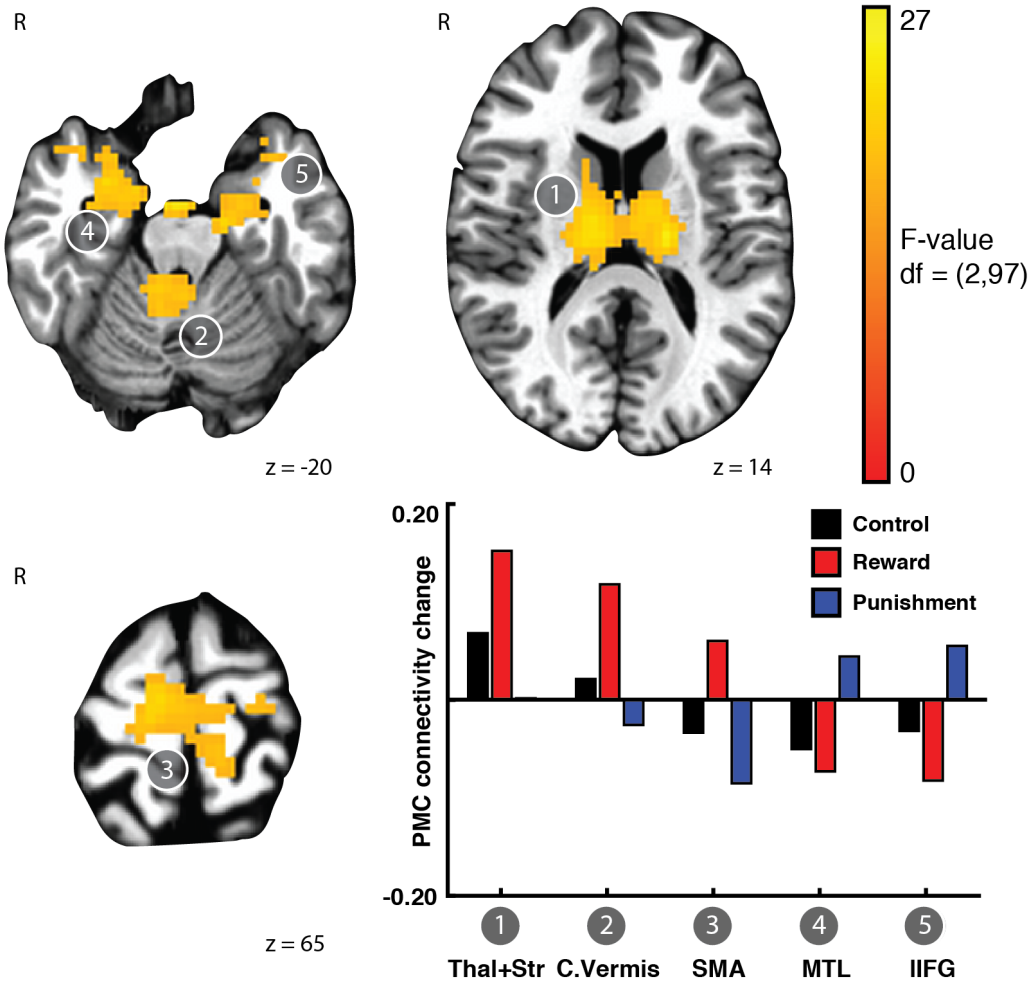


Figure 6.2. Reward and punishment differentially effect PMC functional connectivity change after training on the SRTT. Linear mixed effects model revealed brain regions exhibiting a Rest x Feedback group interaction in the functional connectivity of left PMC. Functional connectivity between PMC and the thalamus and striatum, cerebellum, and SMA increased after training with reward and control feedback but decreased after training with punishment. In contrast, functional connectivity between MTL and left inferior frontal gyrus increased after training with punishment but decreased after training with reward and control feedback. Bars show mean connectivity change after training estimated by the LME model.

6.3.2 Punishment promotes PMC-striatum connectivity after training on the FTT

After training on the FTT, PMC functional connectivity change with bilateral lateral occipital temporal cortex, precuneus, left culmen, cerebellum, right striatum, and right medial frontal gyrus differed across the feedback groups (Figure 6.3). Extracting the estimated mean PMC functional connectivity change of these regions revealed that two primary modes of change in these regions. PMC functional connectivity with the lateral occipital cortex, precuneus, culmen, and cerebellar lobule VI decreased after training with punishment, while PMC functional connectivity with right medial frontal gyrus decreased after training with control feedback. In contrast, PMC functional connectivity with the lateral occipital cortex, precuneus, culmen, and cerebellar lobule VI increased after training with control feedback. In contrast, PMC functional connectivity with right striatum and right medial frontal gyrus increased after training with punishment. Overall, in the regions showing the interaction between functional connectivity change after training and feedback group, reward had a smaller influence compared to control or punishment.

These results suggest that, consistent with the behavioural effects observed at the end of training Chapter 4, compared to reward and control feedback, punishment induces marked changes in functional connectivity change after FTT training while reward and control feedback are more similar.

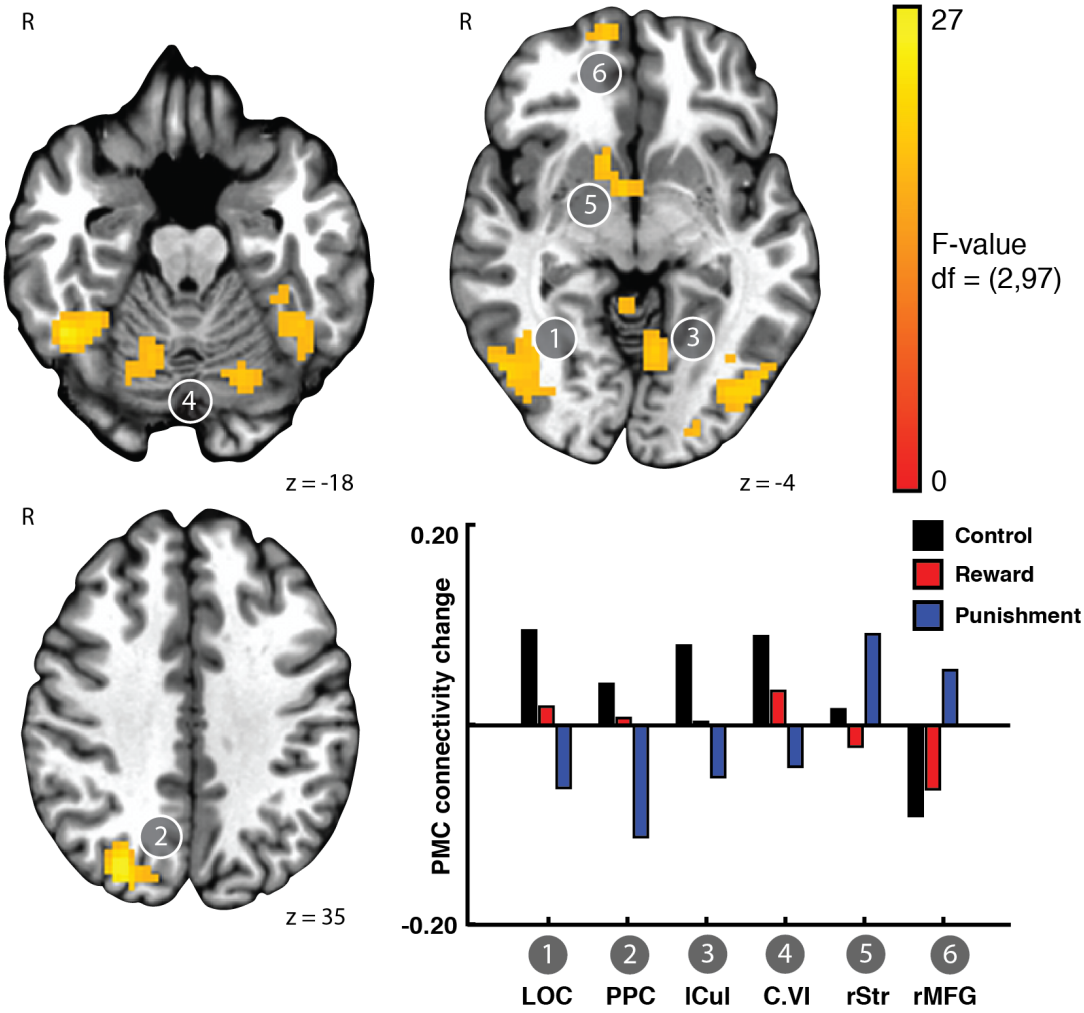


Figure 6.3. Punishment promotes PMC-striatal connectivity after FTT training. The LME model revealed brain regions exhibiting a Rest x Feedback group interaction in the functional connectivity of left PMC. Functional connectivity between the PMC and lateral occipital cortex, posterior parietal cortex, culmen, and cerebellum increased after training with control feedback, but decreased after training with punishment. In contrast, PMC connectivity with right ventral striatum and right medial frontal gyrus increased after training with punishment but decreased after training with reward and control feedback. Bars show mean connectivity change after training estimated by LME model.

Table 6.1. Results of the limited mixed effects model analysis for the SRTT and FTT.

SRTT

Number of voxels	Volume	x	y	z	Region label
Main effect Rest					
1909	51543	-4.5	-22.5	11.5	Left thalamus
949	25623	52.5	1.5	-6.5	Right superior temporal gyrus
477	12879	1.5	-73.5	-24.5	Right cerebellar vermis
299	8073	-55.5	-1.5	-3.5	Left superior temporal gyrus
68	1836	25.5	-13.5	-12.5	Right parahippocampal cortex
Rest x Condition					
828	22356	-16.5	1.5	-0.5	Left pallidum
215	5805	28.5	1.5	-18.5	Right parahippocampal cortex
195	5265	-10.5	-34.5	59.5	Left SMA
154	4158	-4.5	-37.5	-24.5	Left cerebelllum
77	2079	-16.5	-13.5	-18.5	Left parahippocampal cortex
59	1593	-31.5	19.5	15.5	Left inferior frontal gyrus

FTT

Number of voxels	Volume	x	y	z	Region label
Main effect Rest					
3545	95715	-4.5	-22.5	11.5	Left thalamus
148	3996	7.5	-31.5	59.5	Right SMA
118	3186	-1.5	55.5	-6.5	Left medial frontal gyrus
112	3024	43.5	-61.5	-3.5	Left middle occipital gyrus
112	3024	1.5	-61.5	17.5	Right posterior cingulate cortex
60	1620	16.5	-73.5	29.5	Right precuneus
Rest x Condition					
639	17253	46.5	-49.5	-15.5	Right fusiform gyrus
395	10665	-49.5	-64.5	8.5	Left middle temporal gyrus
147	3969	-10.5	-91.5	5.5	Left cuneus
124	3348	25.5	-73.5	38.5	Right precuneus
74	1998	13.5	1.5	5.5	Right pallidum
63	1701	19.5	-61.5	-21.5	Right cerebellum lobule VI
60	1620	-7.5	-64.5	2.5	Left lingual gyrus
55	1485	7.5	61.5	2.5	Right medial frontal gyrus

6.3.3 Influence of reward and punishment on PMC connectivity after training depends on task

In order to better understand the correspondence between the feedback and the change in PMC functional connectivity after training between the two tasks, I calculated the overlap between the sets of regions that showed the Rest x Group interaction for each task separately. Given the reduced power to detect general effects that might be weak but consistent across tasks, for this qualitative assessment I considered the conservative threshold reported above ($p < 0.005$, $k = 54$), as well as a liberal threshold ($p < 0.01$, $k = 54$; Figure 6.4a).

The mean connectivity change with PMC for clusters greater than 20 voxels in extent is shown in Figure 6.4b. At the conservative threshold, one region in the ventral pallidum showed overlap (Cluster 1). When I extracted the estimated mean connectivity change after training for each task, however, I found that the two tasks evoked opposite patterns of connectivity change across the feedback groups. In the SRTT, connectivity between the PMC and the overlapping region increased after training with reward but decreased after training with punishment. In contrast, after training on the FTT, the reward group had increased connectivity between PMC and the overlapping region, and the punishment group had decreased connectivity between PMC and the overlapping region. In both tasks, connectivity between PMC and the overlapping region increased after training with control feedback.

At the liberal threshold, I detected three other clusters in the right anterior temporal lobe, the right cerebellum, and right dorsal thalamus. In each of these clusters, the pattern of connectivity change with PMC exhibited by the reward and punishment groups was also not consistent across tasks.

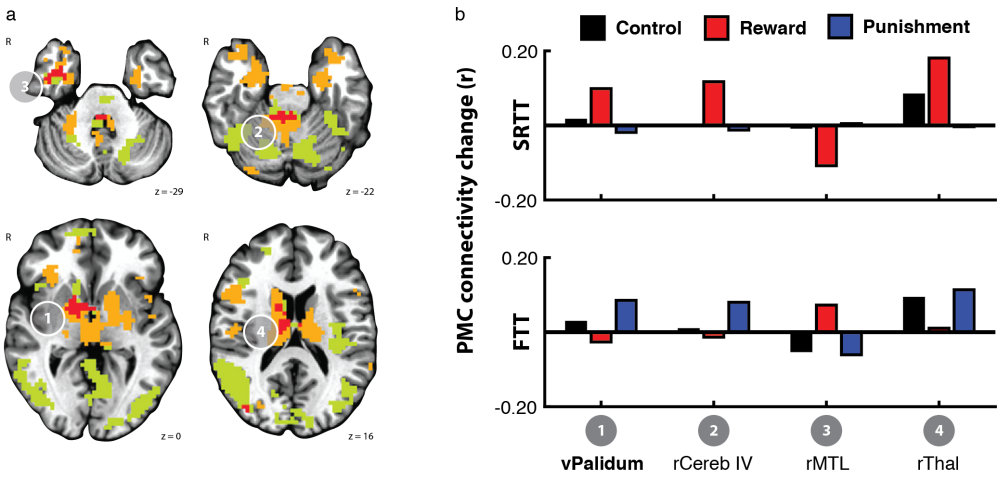


Figure 6.4. Regions showing PMC functional connectivity change after training had limited overlap across tasks. Clusters that showed a PMC connectivity change after training that differed by the feedback group in the SRTT (orange) and the FTT (green) showed limited overlap, even at a liberal cluster threshold ($p < 0.01$, $k = 54$). Only the ventral pallidum (Cluster 1, bold) showed an overlap at the threshold reported above ($p < 0.005$, $k = 54$). In all clusters, the reward and punishment groups exhibited opposite patterns of connectivity change with PMC due to training on the SRTT (upper) and FTT (lower). Bars show mean connectivity change after training estimated by the LME model.

To directly compare the effect of task on the functional connectivity change after training with feedback, I performed a whole-brain analysis (cluster corrected; $p < 0.005$, $k = 54$) using a LME model that included Task (SRTT/FTT), Rest (Pre-/Post-training), and Feedback group (Control/Reward/Punishment) as factors; for this analysis, I specifically focused on regions that showed a Task x Feedback x Rest interaction.

All regions that appeared in the overlap analysis reported above evidenced a Task x Feedback x Rest interaction (bilateral basal ganglia including caudate, putamen, and pallidum [cluster peak: x=-10, y=16, z=17], right cerebellum [1.5, 24.5, -24.5], right MTL [-40, 4.5, -33.5]. In addition, the right lateral temporal lobe [-43, 49, -15], right parietal cortex [-43, 55, 17.5], right medial frontal gyrus [-13.5, 52.5, -6.5], and right lateral prefrontal cortex [-46.5 -16.5 17.5] also showed this interaction (Figure 6.5).

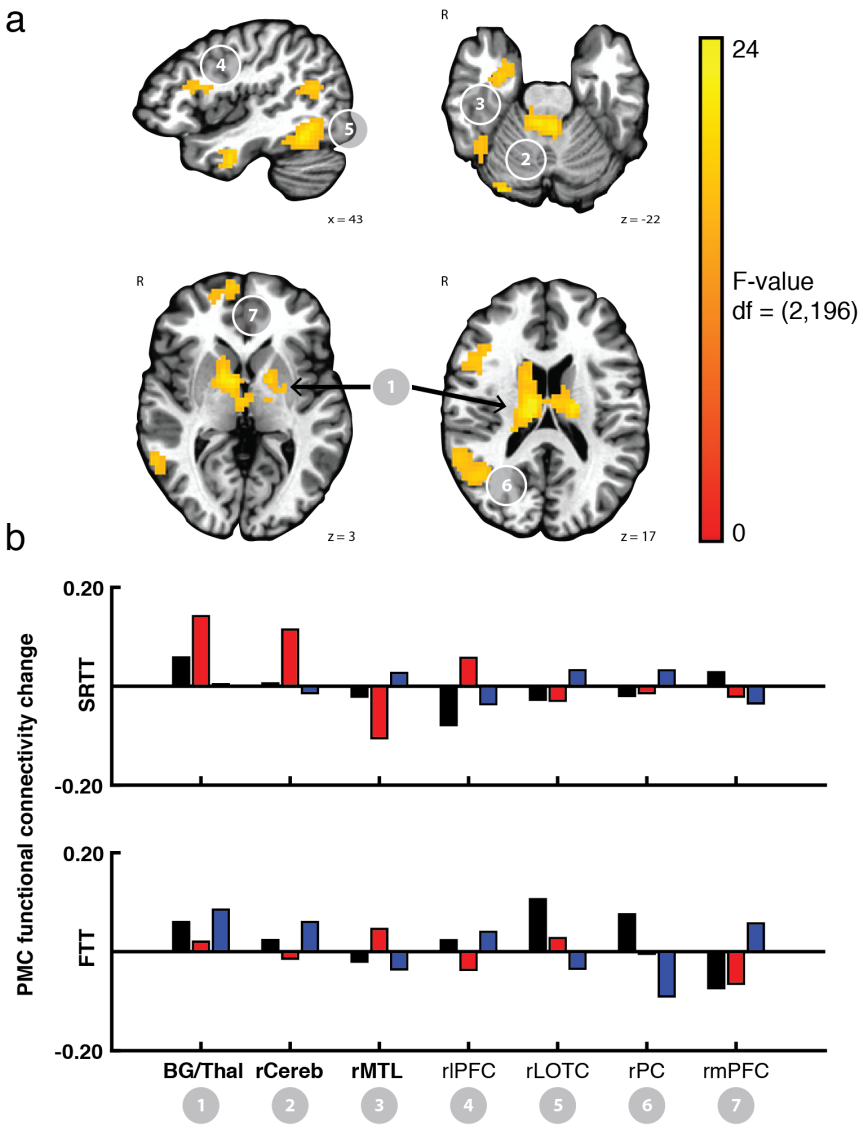


Figure 6.5. Difference in functional connectivity change after training with feedback depends on task. Upper. Data from both tasks were submitted to a LME model. This analysis revealed that the basal ganglia, cerebellum, MTL, temporal cortex, parietal cortex, and prefrontal cortex showed an effect of feedback on functional connectivity change after training that varied by task. Mean connectivity change from these ROIs estimated by the LME is shown in the lower panel.

6.4 Discussion

In this experiment, I sought to investigate the effect of valenced feedback on brain activity during the period immediately following training on the SRTT and FTT. I found that reward and punishment differentially impacted the change in PMC functional connectivity, and this impact was influenced by task. After

training on the SRTT, functional connectivity between PMC with striatum, cerebellum, and SMA increased after training with reward, while functional connectivity between PMC connectivity with MTL and inferior frontal gyrus increased after training with punishment. Training on the FTT induced a different pattern of connectivity change. After training on the FTT, functional connectivity between PMC and parietal cortices increased after training with control feedback but decreased after training with punishment; however, after training on the FTT with punishment, functional connectivity between PMC with the striatum and medial frontal gyrus increased after training. Finally, I looked for regions that showed an overlap in the Rest x Feedback group interaction across tasks. At the conservative threshold, only the ventral pallidum showed an overlap for this interaction, and training the two tasks reward and punishment had opposing effects on connectivity between PMC and this region. After training on the SRTT, connectivity between the PMC and the pallidum increased after training with reward but decreased after training with punishment; the opposite pattern was true after training with the FTT: connectivity between the PMC and the pallidum increased after training with punishment but decreased after training with reward. Taken together, these results suggest that training with valenced feedback has differential effects on functional connectivity after training, and that these effects primarily manifest in the regions involved in task performance.

6.4.1 Differential effects of reward and punishment on PMC connectivity after SRTT training

Reward and punishment differentially affected PMC functional connectivity with the striatum and MTL after training on the SRTT. Others have shown that the SRTT learning engages the motor system, including motor cortex

and PMC [55-57, 313-315], as well as parietal cortex [316], basal ganglia [226, 317-319], and MTL [226, 252, 315, 318]. These structures are often classified into fast (frontoparietal and MTL) and slow (motor and basal ganglia) systems, based on the onset of evoked activity of these regions during the learning process [248], and one theory to explain the impact of reward and punishment on memory formation is the engagement of the slow and fast systems, respectively. With respect to the SRTT, I found evidence for differential recruitment of the fast and slow learning systems after training with reward and punishment. PMC connectivity with motor and basal ganglia structures increased after training with reward, and PMC connectivity with anterior insula and MTL after training with punishment.

In addition to reward and punishment, other factors can interact with memory systems recruited during SRTT learning. For example, the contribution of the motor, parietal, and subcortical systems may be based on the goal of the learner, or based on the statistical complexity of the sequence being learned [157]. However, it is clear that these memory systems interact during the learning process. For example, others have demonstrated that it is possible to bias activity towards one system (e.g., [313]) and to induce competitive interactions among the memory systems [278, 320] by constraining task conditions. There has been evidence to suggest that valenced feedback could also dissociate the memory systems involved during skill-learning [152], and the present data extend this knowledge by showing that reward and punishment can bias recruitment of neural areas during the period after learning.

6.4.2 Punishment influences PMC functional connectivity after FTT training

Reward and punishment also had differential effects on functional connectivity after FTT training, but in different regions compared to the SRTT. After training on the FTT, I observed differences among the feedback groups in functional connectivity between PMC and posterior parietal, lateral occipital, and prefrontal cortices, as well as the ventral striatum and the cerebellum. The control and reward groups showed increases in connectivity between PMC and the parietal, occipital, and cerebellar regions, but decreasing connectivity between PMC and ventral striatal and medial frontal gyrus. The punishment group showed the opposite effect: functional connectivity between PMC and parietal, occipital, and cerebellum decreased after training and PMC connectivity, while the striatum and medial frontal gyrus increased after training.

FTT learning is generally associated with decreased BOLD-signal in the cortical motor network, prefrontal cortex, and caudate nucleus, and increased BOLD-signal in the cerebellum and putamen [15, 321]. While prior studies of FTT training did not find BOLD-signal activity changes in lateral occipital cortex, this region is strongly associated with action representation [322-324] and may be consistently engaged while learning the novel visuomotor association for learning on the FTT. Thus, the effects of feedback on the FTT are largely restricted to areas involved in performance of the task.

6.4.3 The impact of feedback on functional connectivity depends on task

Overall with respect to the influence of feedback on functional connectivity change due to training, I observed very little overlap between the FTT and SRTT.

At the conservative threshold, the only region where feedback modulated functional connectivity change due to training after training on both tasks was the ventral pallidum. In this region, reward and punishment feedback had opposing effects in the two tasks. After training on the SRTT, functional connectivity between the PMC and ventral pallidum increased after training with reward and decreased after training with punishment. The opposite was true after training on the FTT. In participants trained on the FTT, I observed increased functional connectivity between PMC and the ventral pallidum after training with punishment, but a decrease after training with reward. The ventral pallidum contains cell populations that respond to both appetitive and aversive stimuli [325]. Behavioural response to appetitive and aversive stimuli are mediated by glutamatergic and GABAergic neurons in the ventral pallidum, and project to the ventral tegmental area and lateral habenula, respectively [326]. Therefore, given that cells in this area respond to positive and negative stimuli, it may not surprising to find this region affected by both reward and punishment following training in this study.

An explanation for the differential activity of the pallidum across tasks may lie in the behavioural performance. Across both tasks, the dissociation between the feedback groups with respect to PMC-pallidum functional connectivity change after training matched the dissociation behavioural performance: the group where PMC-pallium functional connectivity increased after training performed worse on the task, while the group showing diminished PMC-pallidum functional connectivity performed better at the end of training. Specifically, on the SRTT, the punishment group performed better during training (i.e., had a faster median RT) compared to the reward and control groups, and, on this task, I observed a

decrease in functional connectivity between ventral pallidum and PMC after training with punishment. On the other hand, in the FTT, the punishment group performed worse (i.e., had greater tracking error) than the reward and control groups behaviourally, and the connectivity between PMC and ventral pallidum increased after training with punishment. Thus, with respect to the pallidum, it may be that the connectivity of this region to PMC depends on the experience of the learner (either absence of reward or presence of punishment) in the environment, and this activity may persist after training.

A meta-analysis of decision-making studies using valenced feedback found surprise, valence, and signed prediction error information converged specifically in the ventral pallidum [262]. It is notable that the pallidum itself is a heterogenous structure embedded within the corticobasal-ganglia thalamo-cortical loops [70]. Even though this structure may play a common role integrating the error-related information in both tasks, this error-related information may be originating from distinct cortical networks projecting to separate aspects of the pallidum (e.g., associative versus motor) [70]. The scanning resolution used here does not permit localizing specific areas within the pallidum, but this could be an interesting aspect for follow-up with higher-resolution imaging modalities.

Even at a liberal threshold, the regions that showed overlap across the two tasks showed opposite patterns of connectivity change based on the feedback given. Similar to the pallidum, the differential effect of feedback on functional connectivity between PMC and these regions may be due to the qualitative differences in performance. More generally, the minimal overlap and opposing patterns suggest that reward and punishment strongly impact post-training brain activity, and that this effect is strongly influenced by the task being learned.

6.4.4 Limitations

Several aspects of the present study are worth noting.

Firstly, we did not distinguish between dorsal and ventral PMC in the present study because no regions showed a Rest x Group x ROI interaction in either task. Dorsal and ventral PMC are highly interconnected and although feedback may differentially impact dorsal and ventral PMC, we may have been underpowered to detect any effect.

Secondly, my implementation of the SRTT included a fixed trial length rather than a self-paced trial duration, which might foster explicit knowledge, and it has been shown that explicit learning recruits different neural networks after learning [150]. No participants included in this experiment spontaneously reported sequence knowledge. When tested at 3+ weeks, participants showed no evidence of explicit awareness (for further discussion, see Chapter 4). In addition, in order to match the SRTT parameters, the FTT implementation used here featured continuous feedback of cursor position, rather than feedback at given at the end of the trial, and future work may consider the difference between continuous versus end-of-trial feedback.

Thirdly, I did not find any differences in retention across the feedback groups in either task, despite observing differences in the acquisition and performance during training. Therefore, I cannot say whether any specific PMC connections facilitate better memory formation.

Fourthly, it is widely known that sleep interacts with memory formation and offline memory processes [226, 277, 327, 328]. In this experiment, I did not measure sleep quality. Whether sleep interacts with the effect of reward and punishment on retention may be an interesting avenue for future research.

Fifthly, because of scan time restrictions, I could not precisely control for the time-of-day that participants were scanned as presented in Chapter 3. In this experiment, I controlled for time of day by randomizing the subject assignment to ensure that there was no systematic bias across the study groups. The extent to which time-of-day impacts reward system function is an area of ongoing research [329, 330] and may be of interest for future work, as well.

Finally, because the two tasks are not inherently comparable on the basis of behaviour, the number of participants (n=12) was severely underpowered for the purpose of performing correlational analyses. Therefore, I elected not to perform these analyses and focus on the change in correlation after training. I discus this issue in greater detail in Chapter 8.

6.5 Conclusions

In summary, over the course of two behavioural experiments and one imaging experiment, I have shown that:

1) The impact of valenced feedback, and punishment in particular, depends on the task demands.
2) Even when performing a single task, the goal of the learner modulates the effect of valenced feedback.
3) Consistent with this behavioural data, task modulates the influence of feedback on the cortical and subcortical areas engaged in the period immediate after learning.

As described in the Chapter 1, because of the brain's hierarchical organization, skills may be acquired along multiple domains (e.g., motor or perceptual) depending on the current demands. Because different brain systems

may be better suited to particular kinds of information, these may become engaged under certain circumstances. The three studies here show that reward and punishment may interact with the way a task is learned and the brain activity evoked during learning.

Collectively, these results provide a fuller picture of the adaptability of behavioural performance, highlight the brain's flexibility for learning different skills, and underscore the importance of viewing functional plasticity change as a property of the whole organ rather than in any particular brain region.

In both the present Chapter and Chapter 3, I investigated the functional plasticity using functional connectivity measured using fMRI. However, the underlying physiology that gives rise to functional connectivity is not well understood. In the next Chapter, I use a novel MRSI technique to address this knowledge gap by linking functional connectivity measurements to neurochemical distributions across the brain.

7 Towards a link between network-level neurochemical concentration, functional connectivity and sensorimotor behaviour

In Chapters 3 and 6, I explored the neural consequences of motor skill learning in different contexts, focusing on the impact of learning on the co-fluctuation of the BOLD-signal between brain areas measured with fMRI. However, fMRI is inherently limited, as it provides only an indirect measure of the neurochemical substrates that underlie brain function. Magnetic resonance spectroscopic imaging, a tool that allows for measurement of neurochemicals across a wide area of the human brain, is well-suited to address this limitation. In this Chapter, I implement magnetic resonance spectroscopic imaging for the first time in sensorimotor research, with the goal of linking neurochemicals, network-level functional connectivity, and sensorimotor behaviours.

7.1 Introduction

As discussed in Chapter 2, the BOLD-signal measured with fMRI (used in Chapters 3 and 7) can provide great insight into brain function, but this signal only provides an indirect measurement of the electrical and chemical signals that actually underlie neural processing. For example, somatotopic mapping using fMRI [225, 331], can provide a high-resolution map of where a location on the body is represented within primary somatosensory cortex, but fMRI alone cannot explain how these maps are formed and maintained. Likewise, resting-state functional connectivity networks can be observed reliably using fMRI, and the connectivity strength in these networks can be linked to behaviour [332, 333], but fMRI cannot be used to ascertain the underlying biological determinant of network-level connectivity strength. In order to address these questions, one must measure the chemical substrates for neural function.

One means to measure chemical compounds in the living human brain is MRS. Using MRS, neurochemicals such as GABA and glutamate, the brain's primary inhibitory and excitatory neurotransmitters, respectively, can be quantified. MRS measured GABA levels, which are thought to reflect local inhibitory tone [224], have been linked to many behavioural traits. For example, GABA concentration in visual cortex has been linked to perceptual acuity and altered GABA levels are correlated with perceptual deficits in neurological disorders [334-336]. In the somatosensory and motor systems, GABA levels have also been linked to tactile sensory acuity [225], RT [224], and sequence learning ability [224].

MRS-measured GABA is also related to fMRI measurements. Notably, MRS-measured GABA has been linked to the tuning of somatotopic maps in primary

sensory cortex [225]. In addition, motor cortical GABA-concentration is related to the motor-network connectivity strength [153], and GABA concentration in medial parietal cortex is related to default mode network connectivity strength [154].

The studies described above have provided significant insights into the role of local inhibition in cortical function and behaviour. However, they have been intrinsically limited as they were only able to quantify neurochemicals from a single MRS voxel. Specifically, in the context of sensorimotor research, the majority of papers using svMRS have placed the voxel over the hand area of the precentral and postcentral gyri. However, as discussed in Chapter 1, multiple brain areas, including medial prefrontal cortex, lateral prefrontal cortex, and SMA also contribute to sensorimotor tasks. In addition, as described in Chapter 1, 3, and 6, the interaction between multiple brain regions is potentially crucial for performance and learning of skills. Therefore, it is necessary to extend neurochemical investigations to multiple voxels in order to understand how the distribution of neurochemicals across the brain facilitates covariation in brain activity.

In the present Chapter, I use the novel MRSI technique described in Chapter 2 to begin investigating these questions. I had three primary goals. First, I sought to demonstrate the reliability of this novel MRSI acquisition technique for measuring GABA and glutamate (together with glutamine; Glx). Second, I wanted to develop a strategy suitable for quantifying regional neurochemicals so that they may be related to individual differences in sensorimotor ability. Finally, I wanted to investigate a link between neurochemical concentration, resting-state functional connectivity, and sensorimotor performance.

7.2 Methods

7.2.1 Procedure

The present study was conducted across two sessions (Figure 7.1). In the first session, I collected multimodal MRI data, which included a high-resolution T1-weighted scan, a resting-state fMRI scan, a diffusion-weighted imaging scan (not discussed here), and a MRSI scan. In the second session, participants completed the SRTT, action selection task (AST), and the temporal order judgement task (TOJT). This enabled me to characterize the relationship between a broad spectrum of physiological and sensorimotor-performance traits at the individual subject-level.

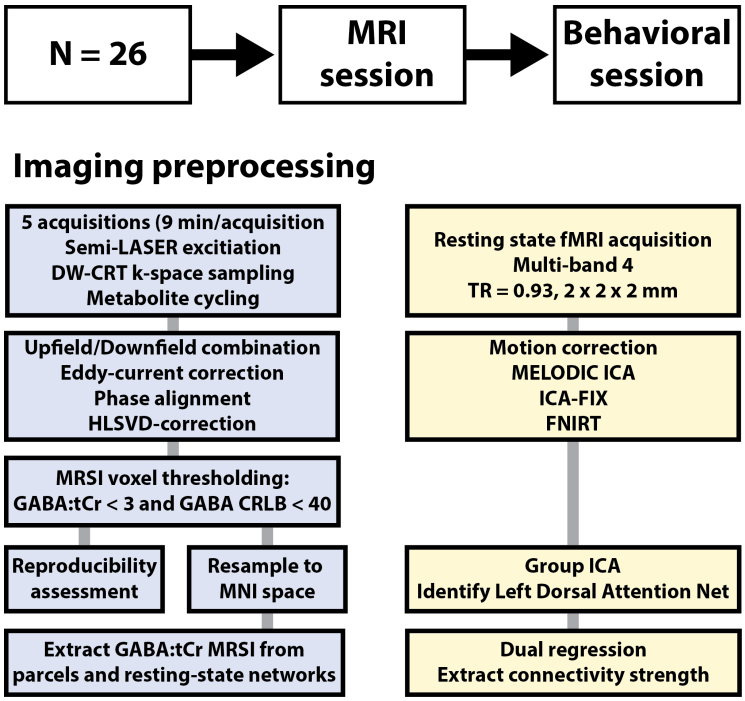


Figure 7.1. Experimental design. Top. The experiment took place over 2 sessions. In the first session, participants underwent a multimodal MRI battery, including MRSI and resting-state fMRI. In the second session, participants performed a series of behavioral tests. Middle. The processing pipelines for both MRSI (blue) and resting-state fMRI (yellow).

7.2.2 Participants

Twenty-six (16 female, 23.9 ± 2.6 years old) right-handed participants were recruited and gave their informed consent in accordance with central university research ethics committee regulations (MSD-IDREC-C1-2014-100(R446/RE002)).

Seven participants failed to perform the AST properly due to misunderstanding the instructions (see Section 7.3.5.2). Rather than perform Execute blocks, three participants continued to perform the Select blocks. In addition, four participants failed to follow the appropriate button pressing assignments during Select blocks. These participants were therefore discarded from the analyses relating motor performance to behaviour.

7.3 MRI acquisition

All data was collected on a Siemens MAGNETOM Prisma 3T MRI scanner (Siemens, Erlangen, Germany).

7.3.1 T1-images

7.3.1.1 T1 acquisition

For registration purposes, a T1-weighted anatomical image was acquired with the following parameters (magnetization-prepared rapid gradient echo (MPRAGE), TR = 1900 ms, TE = 3.96 ms, TI = 912 ms, flip-angle = 8 degrees, bandwidth = 51.2 kHz, FOV = 256x232 mm, acquisition matrix = 256 x 232, resolution = 1 x 1 x 1 mm, 192 slices per volume).

7.3.1.2 T1-based tissue quantification

Cortical parcellations for each subject were derived using Freesurfer [236]. Voxel-wise grey matter probability determined using FAST 4.0 implemented in FSL [337] served as a proxy for grey matter density for MRSI analyses using the canonical MRS preprocessing described below.

7.3.2 Magnetic resonance spectroscopic images

Ten MRSI scans (time per scan = 4.5 minutes, 45 total minutes) were collected using the non-water suppressed semi-LASER DW-CRT sequence presented in [204]. A full description of the acquisition parameters for this scan is given in Section 2.3 (Methods: Magnetic resonance spectroscopic imaging).

7.3.2.1 Density-weighted concentric ring trajectory

The DW acquisition method relies on varying the density of sampled points in k-space to create an effective windowing of the data. A Hanning-window density weighting was used to determine the radius of each ring, since it has been shown to provide a good compromise between sensitivity and spatial resolution. The function used is given by:

$$\text{Eq. 1} \quad w(k) = \frac{\Delta x}{2}(1 + cos(2\pi k \Delta x/\alpha))$$

where Δx is the nominal spatial resolution. In this equation, the parameter α controls the adjustment of the effective spatial resolution by modulating the width of the SRF main lobe. When α is set to 1, DW-CRT spreads between $-1/(2\Delta x)$ and $1/(2\Delta x)$.

7.3.2.2 Metabolite cycling

During MRSI acquisition, participants viewed a nature documentary ('Mountains', *Planet Earth*, BBC) with closed captioning on. To ensure they did not

fall asleep during the scan, the participant was told to indicate with a button press when any word describing 'mountains' appeared on screen (<1 button press per minute).

Two asymmetric narrow transition-band adiabatic RF pulses with mirrored inversion profiles were applied in alternating TRs for the inversion of the upfield and downfield (relative to water) spectral resonances 9.6 ms before the semi-LASER localization [209]. An 80 Hz transition bandwidth ($-0.95 < M_z/M_0 < 0.95$) and 820 Hz inversion bandwidth ($-1 < M_z/M_0 < -0.95$), 70 to -750 Hz) downfield/upfield from the carrier frequency was achieved using a maximum B_1 of 19 μT and pulse duration (T_P) of 27 ms. The center of the transition band ($M_z = 0$) was placed at the carrier frequency offset by +60 Hz and -60 Hz for downfield and upfield, respectively. The pulse sequence was delivered with adiabatic pulse parameters: hyperbolic secant pulse, $HS_{1/2}$-, with R = 10 and 0.9 x T_p, tanh/tan pulse with R = 40 and 0.1 x T_P).

2D MRSI scans were collected from a single slice manually positioned using the T1-weighted MP-RAGE image. The slab was placed as dorsal as possible to the cingulate gyrus while being contained entirely within brain tissue. Regions targeted for coverage included motor cortex and medial and lateral prefrontal cortex, given their important role in sensorimotor performance described in Chapter 1.

A 2D concentric *k*-space trajectory was used to sample polar *k*-space data [214]. To satisfy the requirement of avoiding azimuthal aliasing [338], 64 points per ring were collected with an ADC bandwidth of 80 kHz and a maximum slew rate = 138.5 mT/m/ms. 512 spectral points (N_{sp}) were collected with an effective spectral bandwidth of 1250 Hz. Thus, a high in-plane resolution (5 mm x 5 mm)

with thickness 15 mm was achieved using FOV = 240 mm x 240 mm x 10 mm, semi-LASER localization = 80 mm x 75 mm x 10 mm, TR = 1400 ms, TE = 32 ms, 4 spatial interleaves ($N_{interleaves}$), $N_{directions}$ = 2 (upfield/downfield), N_{ring} = 24, (TR x $N_{interleaves}$ x N_{ring} x $N_{directions}$ = 268.8 s).

7.3.2.3 MRSI preprocessing

Reconstruction algorithms were implemented in MATLAB (MathWorks, Natick, MA, USA). Non-uniform fast Fourier transform (NUFFT) gridding was performed without using any post-hoc density compensation, since this is not required for DW-CRT data. The final matrix size of the reconstructed MRSI image after NUFFT was $2N_{ring}$ x $2N_{ring}$ x N_{sp} x $N_{channels}$ x $N_{directions}$.

Odd (upfield) and even (downfield) single-shot non-water-suppressed FIDs edited by asymmetric RF pulses were corrected on the basis of the water signal. The frequency correction was performed using a cross-correlation algorithm and phase correction was performed using a least squares fit algorithm. Upfield and downfield edited spectra were summed. The combined upfield and downfield spectra were then used to remove residual eddy current effects [339] and combine the phased-array coil spectra [340]. Metabolite spectra were subsequently calculated by subtracting the alternating FIDs. Differences in the water peak amplitude between upfield and downfield edited spectra was corrected using the method reported in [341]. The residual water peak was filtered with the HLSVD algorithm [218].

Prior to LCModel fitting, two consecutive spectra (i.e., image 1-2, image 3-4, etc.) were averaged together to increase spectral signal-to-noise prior to LCModel fitting. Thus, each participant had 5 images for their final image set.

7.3.2.4 Voxel-wise metabolite quantification

The LCModel package was used to quantify the metabolite spectrum for each MRSI voxel [220]. The model spectra of alanine, aspartate, ascorbate/vitamin C, glycerophosphocholine, phosphocholine (PC), creatine (Cr), phosphocreatine (PCr), GABA, glucose, glutamine (Gln), glutamate (Glu), glutathione, lactate (Lac), myo-Inositol (myo-Ins), NAA, N-acetylaspartylglutamate, phosphoethanolamine (PE), scyllo-Inositol and taurine were generated based on previously reported chemical shifts and coupling constants by the GAMMA/PyGAMMA simulation library of VeSPA (Versatile Simulation, Pulses and Analysis) according to a density matrix formalism [342]. Simulations were performed using the same RF pulses and sequence timings as those on the 3T system in use. Eight LCModel-simulated macromolecule resonances were included in the analysis at the following positions: 0.91, 1.21, 1.43, 1.67, 1.95, 2.08, 2.25 and 3 ppm. If the correlation between two metabolites was consistently strong (correlation coefficient < −0.5) in a given region, their sums are reported, e.g., combined glutamate and glutamine (Glu+Gln, Glx), total creatine (Cr + PCr, tCr) and total choline (GPC + PC, tCho). Metabolite data are reported relative to tCr.

All analyses, including analyses of MRSI voxel-wise reliability, only included voxels with GABA:tCr values lower than 0.375 and GABA CLRB values less than 40%.

7.3.3 MRSI alignment to T1 image

Following LCModel fitting, metabolite maps were registered to the participant's T1-weighted image and resampled to 1 x 1 x 1 mm using a nearest-neighbour interpolation.

To relate the participant's metabolite measurements to resting-state functional connectivity measurements, the MRSI derived metabolite maps were warped to MNI space and resampled using nearest neighbour interpolation. To facilitate comparison of regional neurochemicals to behaviour the average metabolite concentrations within each parcel from the Desikan atlas [343] were calculated for each participant. Figure 7.2 shows the alignment of a representative participant's MRSI slab with the Left attention network and the Desikan atlas in standard space. Only parcels/networks with adequate MRSI overlap (defined below) in all participants were considered for subsequent analyses.

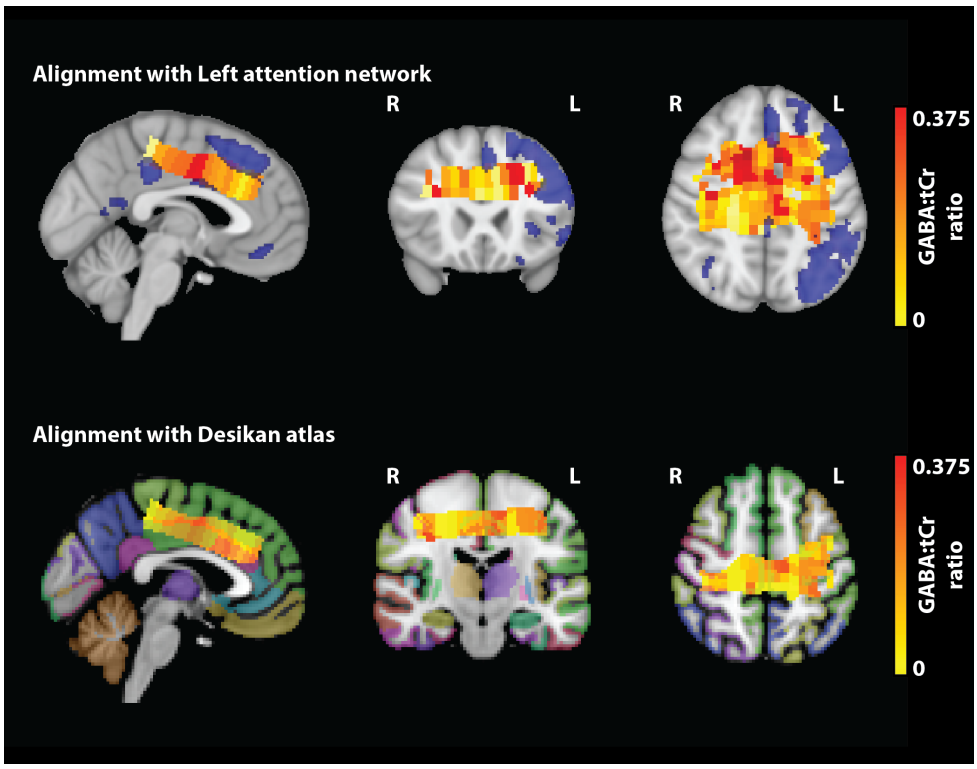


Figure 7.2. Example participant's aligned MRSI slabs. (Top) A representative MRSI slab aligned to the group-ICA defined Left attention network (overlaid in transparent blue). (Bottom) The same participant's slab aligned to the Desikan atlas. Desikan atlas parcels are restricted to grey matter and colored to show distinct boundaries. Only parcels/networks with adequate overlap in all participants were considered for additional analyses.

7.3.4 Resting-state fMRI

7.3.4.1 Resting-state fMRI acquisition

Prior to MRSI acquisition, I collected 7 minutes of resting-state fMRI. During the resting-state scan, participants were instructed to fixate on a white cross projected in the centre of a black background. The scan was acquired using the following parameters: TR = 933 ms, TE = 33.4 ms, multiband acceleration factor = 6, flip angle = 60 degrees, bandwidth = 43.4 kHz, FOV = 192 x 192 mm, acquisition matrix = 96 x 96, resolution = 2 x 2 x 2 mm, 72 slices per volume, number of volumes = 451.

7.3.4.2 Resting-state fMRI preprocessing

Resting-state fMRI data were preprocessed using FSL [199]. Brain extraction on structural T1 images was performed using BET and the anatomical image was subsequently aligned to the MNI-152 template using non-linear boundary-based registration with FNIRT. For EPI data, motion correction was applied using MCFLIRT, followed by distortion correction (B0 unwarping), and high-pass filtering. These data were subsequently registered to the T1-weighted image. For denoising, the registered EPI data were decomposed into independent components using MELODIC. For each subject, the components were manually classified as signal or noise. The time-series from noise components were then regressed out of the data (*fsl_regfilt*). The denoised data were then warped to standard space. No additional smoothing was applied to the resting-state data.

7.3.5 Behavioural tasks

Procedures for the SRTT, AST, and TOJT are described below (Figure 7.3). The order of the tasks was counterbalanced across participants. All tasks were implemented using PsychoPy 2.8 [234].

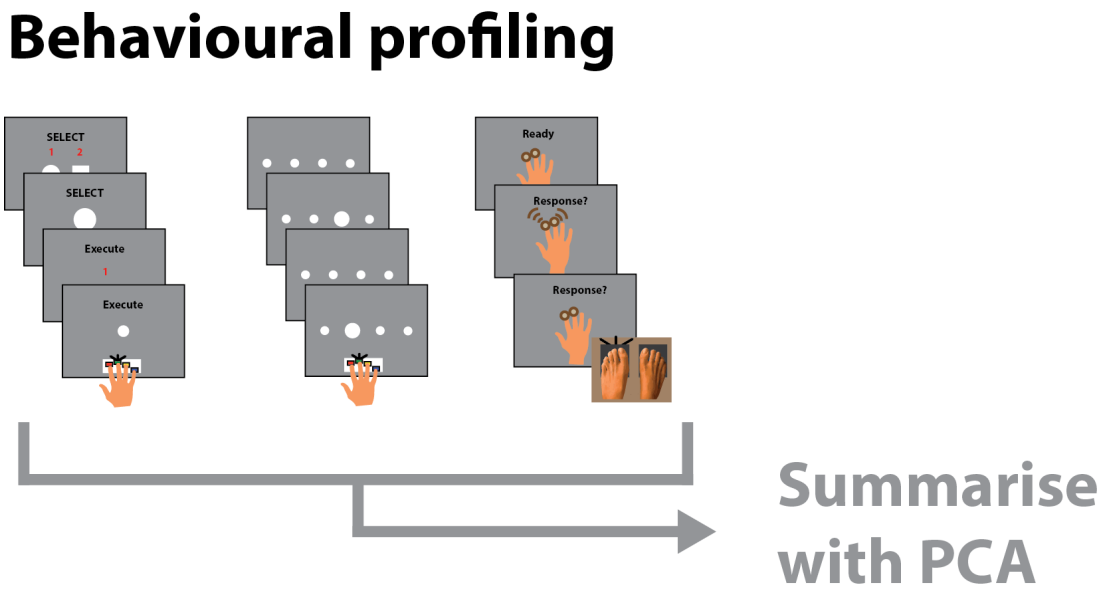


Figure 7.3. Three sensorimotor tasks were performed to define the sensorimotor score. Participants completed the Action Selection Task (AST), the Serial Reaction Time Task (SRTT), and the Temporal Order Judgement Task (TOJT). The performance outcomes for each (AST: Action selection proficiency, SRTT: Sequence learning and Chunking, TOJT: Perceptual acuity) were summarized using PCA to produce a Sensorimotor performance score.

7.3.5.1 Serial reaction time task

For the SRTT, participants were seated in front of a computer screen and held a button box with their right hand. At the beginning of a trial, four small circles appeared on screen. When a large circle appeared in a one of the positions, the participant pressed the corresponding button. After a button press, the stimulus would disappear. The next stimulus would appear after 150 ms. Participants were told to respond as fast and accurately as possible.

Trials were grouped into blocks of 64 trials. During Blocks 1, 6, 11, 16, and 18, the stimulus would appear in random positions (but appearing no more than two times in consecutive locations). During all other blocks, the stimulus would appear according to a fixed sequence (16-items long). Participants were randomly assigned to one of two possible sequences: 1) [2,2,4,3,3,2,1,1,4,2,3,4,4,1,3,1], or 2) [3,3,4,3,1,1,2,4,4,1,3,2,2,1,4,2]. Participants were not informed about the presence of the sequence. Between blocks, a cross would appear on screen. After a mandatory 2 s break, the participant could press the spacebar to move on.

For both sequence and random blocks, inaccurate trials and trials with RTs below 100 ms or greater than 800 ms were excluded from the analyses. Sequence learning was calculated by comparing the median RT during the random blocks to the median RT during the sequence blocks (early, middle, and late probes). Sequence knowledge, the summary score for this task, was equal to the difference between sequence learning scores calculated from early and late probes.

In addition to sequence knowledge, chunking scores were calculated for each subject. Chunking captures the extent to which participants subdivide the sequence into smaller segments. Chunk boundaries can be delineated based on RT: RTs within a chunk will be faster than those at the boundaries of a movement chunk. As a person's chunking structure becomes more defined, the RTs at each position of a sequence will become more defined. In contrast, participants concatenating chunks will show less consistent RTs at each position of the sequence. For each sequence block, I calculated the median RTs for each position in the 16-item sequence. The median RT at each position in the 16-item sequence was then correlated with the other sequence blocks. The difference between the mean correlation of the early (Blocks 2-5) and the late (Blocks 12-15) periods

yielded the chunking score, which differentiates participant's ability to develop chunks early and late in learning.

7.3.5.2 Action selection task

The AST was implemented in the same manner as [52, 344]. Participants were seated in front of a computer screen with a button box in the right hand. For a trial, one of four shapes would appear on the screen (large/small circle/square). Trials were grouped into two types of blocks, 'select' and 'execute'. For select blocks, participants responded with the index or middle finger to the appearance of the shape using a prescribed rule, which was instructed at the onset of the block. The two rules were: A) large circle or small square instructs index finger, small circle or large square instructs middle finger; or B) small circle or large square instructs index finger, large circle or small square instructs middle finger. In execute blocks, participants responded as fast as possible to any shape using the finger instructed at the beginning of the block. Blocks consisted of 48 trials, and participants completed 3 blocks of each type. Inaccurate trials were excluded from the analysis. Action selection proficiency was calculated by comparing the median RTs between the select and execute blocks, normalized by the median RT during execute blocks.

7.3.5.3 Temporal order judgement task

The TOJT was implemented using the same procedure as [225]. Participants were seated in front of a computer with the right index and middle finger placed on the a vibrotactile stimulator. For a trial, vibrations were delivered to the right or left finger offset by a certain inter-pulse interval. Participants were required to indicate the finger that was stimulated first by pressing a pedal with

the corresponding foot. The temporal offsets were (in ms): [300, 150, 100, 70, 50, 30, 15]. Participants had to indicate their choice within 1.5 s.

A psychometric function was fitted to the accuracy at each temporal offset. Sensory acuity was defined as half the difference between the two temporal offsets that give a "middle finger first" judgement in 75% and 25% of the trials (just noticeable difference, JND). Smaller values of JND reflect better perceptual acuity.

7.3.5.4 Sensorimotor summary score

These three tasks yielded the following scores: Action selection score (from the AST), Motor learning (SRTT), Chunking (SRTT), and Perceptual acuity (TOJT). These scores were independently related to the neurochemical measurements. In addition, the participants score on the four measures were submitted to principal components analysis (PCA) to establish an overall "Sensorimotor score" to summarize the data across these measures. I used the Sensorimotor score as my primary metric to link to the neurochemical and resting-state functional connectivity measurements.

7.3.6 MRI analysis

7.3.6.1 MRSI GABA reliability

The first goal of this project was the estimate the reliability of GABA and glutamate measurements using this MRSI approach. To this end, for each participant, I calculated two metrics. First, to assess the reproducibility of the image measured across repetitions, I calculated the similarity (Pearson's r) between each of the GABA and Glx images acquired from the 5 imaging runs. Second, to assess the variability of the measurement at each voxel, I calculated the

voxel-wise coefficient of variability (metabolite standard deviation / mean) across the 5 metabolite images for GABA and glutamate. For comparison, the same procedure was conducted for tCho and tNAA.

These metrics provide an estimation of the reliability of the metabolite measurements and also provide a theoretical 'noise ceiling' for any relationship with behaviour [345].

7.3.6.2 ROI concentration assessment

To quantify GABA and Glx, I calculated the mean values from all non-zero voxels from the resampled MRSI voxels in each Desikan-atlas parcel with adequate coverage in each participant. Adequate coverage was defined as all participants having at least 50 voxels present in a parcel. This left eight parcels to be considered in the analysis. The median coverage (± interquartile range) of each parcel (in mm^3) is listed here: left middle frontal gyrus = 173 ± 75; left paracentral gyrus = 200 ± 59 ; left posterior cingulate = 269 ± 136; left precentral gyrus = 135 ± 74; left superior frontal gyrus = 450 ± 206; right middle frontal gyrus = 92 ± 85; right posterior cingulate cortex = 348 ± 163; right superior frontal gyrus = 317 ± 285.

7.3.6.3 Nuisance correction

Prior work has shown that the importance of correcting for grey matter concentration and other nuisance factors when examining correlations between neurochemical concentration and behaviour (e.g., [120]). Here, I explored two different strategies for correcting for nuisance factors. In the first correction strategy, I used a more canonical approach similar to corrections applied for group-level analyses in other imaging modalities (e.g., [175]): across participants,

I regressed out the nuisance variables LCModel SNR, participant age, and grey matter concentration for each MRSI parcel.

The second correction strategy was novel. For this strategy, I assumed that the uncertainty in the neurochemical estimation would result from (and therefore be indicative of), potential confounding variables (including age, SNR, and grey matter concentration). Given that the CRLB values give the uncertainty (relative to the mean of the concentration), I regressed these out of the data as the only nuisance variable. No other variables were regressed out of the data for the CLRB correction.

Both of the correction strategies above were applied at the group-level (i.e., across participants) and yielded corrected GABA and Glx concentration for each participant centred on 0. After the two different correction strategies were applied, the resulting corrected concentrations were highly correlated for both GABA and Glx (Figure 7.4), suggesting that they remove similar variance from the initial data.

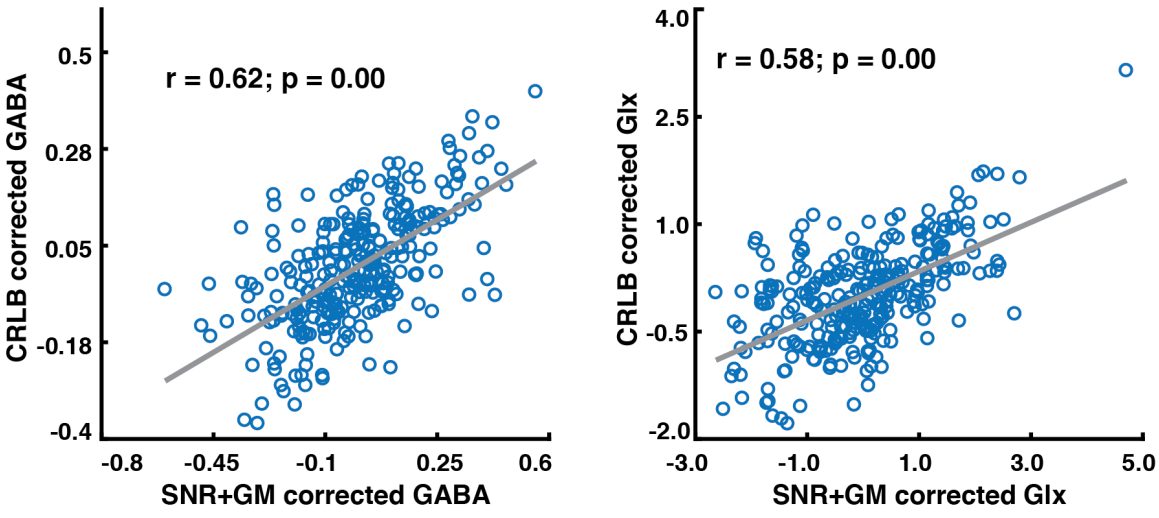


Figure 7.4. Correlation between the SNR+GM correction (regressing out the effects of SNR, grey matter concentration, and participant age) and the novel Cramer-Rao Lower Bound (CLRB) correction for both GABA (left) and Glx (right). The two correction strategies yield highly correlated results. Each point represents the corrected neurochemical value from a single Desikan-atlas parcel for a single participant.

7.3.6.4 Association between neurochemical concentration and behaviour

To assess the relationship between regional neurochemical concentration and behaviour, I correlated the regional GABA concentrations with the Sensorimotor score corrected for the number of regions tested (corrected $p = 0.05/8 = 0.0063$).

For completeness, I also tested for correlation between regional neurochemical concentrations and the four behavioural metrics from the tasks performed (Action selection proficiency, SRTT Learning, SRTT Chunking, and Perceptual acuity).

7.3.6.5 Relationship between resting-state functional connectivity and neurochemical concentration

First, in order to establish a measure of resting-state functional connectivity to relate to the neurochemical concentration, a group-ICA was performed to establish a set of resting-state networks. Fifty independent components were considered, and 16 were determined to reflect BOLD-activity (rather than noise) based on their spatial and temporal characteristics.

The functional connectivity strength for each participant was then calculated using the dual-regression procedure [203, 332]. Briefly, the group-level spatial independent components were regressed back into each participant's resting-state data to establish a time course of each component for each participant. These participant-specific time-courses were then fit to each participant with a GLM. To establish functional connectivity network strength, the mean beta-value of the fit component time-series was computed using the original spatial independent component as a mask (original z-threshold = 2.7). This gave a

single-value for connectivity strength for each network for each participant. The participant's age was regressed out as a nuisance covariate prior to the analysis relating network-level functional connectivity strength to neurochemical concentration.

To relate network-connectivity strength to the neurochemical concentration in each network, the same procedure for calculating the metabolite concentration for the Desikan atlas parcels was applied to the resting-state networks. Because of the substantially larger size compared to the Desikan atlas ROIs, only resting-state networks overlapping with at least 100 resampled MRSI voxels were considered. This criterion left 4 networks to be considered: left attention network (overlapping voxels ± interquartile range = 404 ± 209), dorsal sensorimotor network (426 ± 277), = ventral sensorimotor/auditory network (613 ± 144), cinguloopercular network (827 ± 826.5). The corrected GABA concentration was then related to the network-level connectivity strength of each network using Pearson's correlation.

7.4 Results

7.4.1 DW-CRT enables highly-reproducible GABA measurements at 3-Tesla

Using the novel MRSI acquisition and preprocessing strategy described in the Chapter 2, I was able to acquire extremely high-resolution metabolite measurements from all 26 participants scanned. Five individual participant maps for the metabolites GABA, Glx, tCho, and tNAA (without masking applied) are shown in Figure 7.5. The corresponding CRLB maps are shown in Figure 7.6, and voxel-wise SNR and LW maps for these participants are shown in Figure 7.7. In

addition, for a single participant, the metabolite map from each imaging run, as well as nine representative spectra are shown in Figure 7.8. The median and interquartile range of the SNR, LW, and CRLB values for GABA, Glx, tCho, and NAA are shown in Table 7.1.

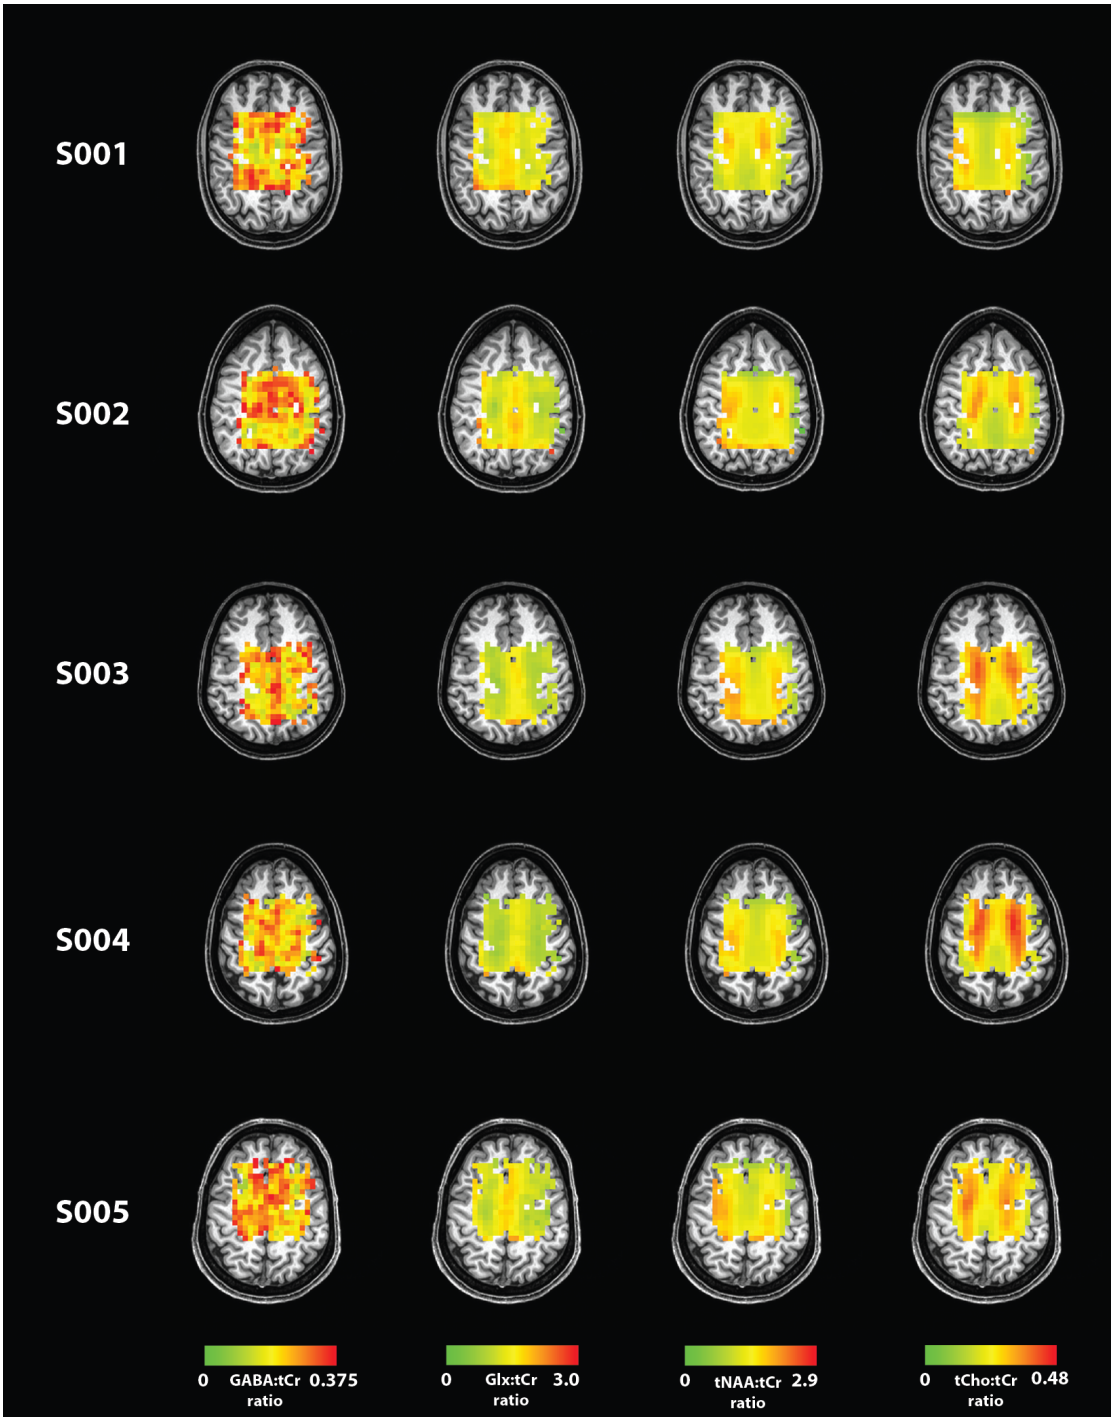


Figure 7.5. Metabolite distribution for GABA, Glx, tNAA, and tCho for 5 example participants. As predicted, Glx and GABA appear in relatively high-concentrations in grey matter, while tCho and tNAA appear in relatively high-concentrations in white matter.

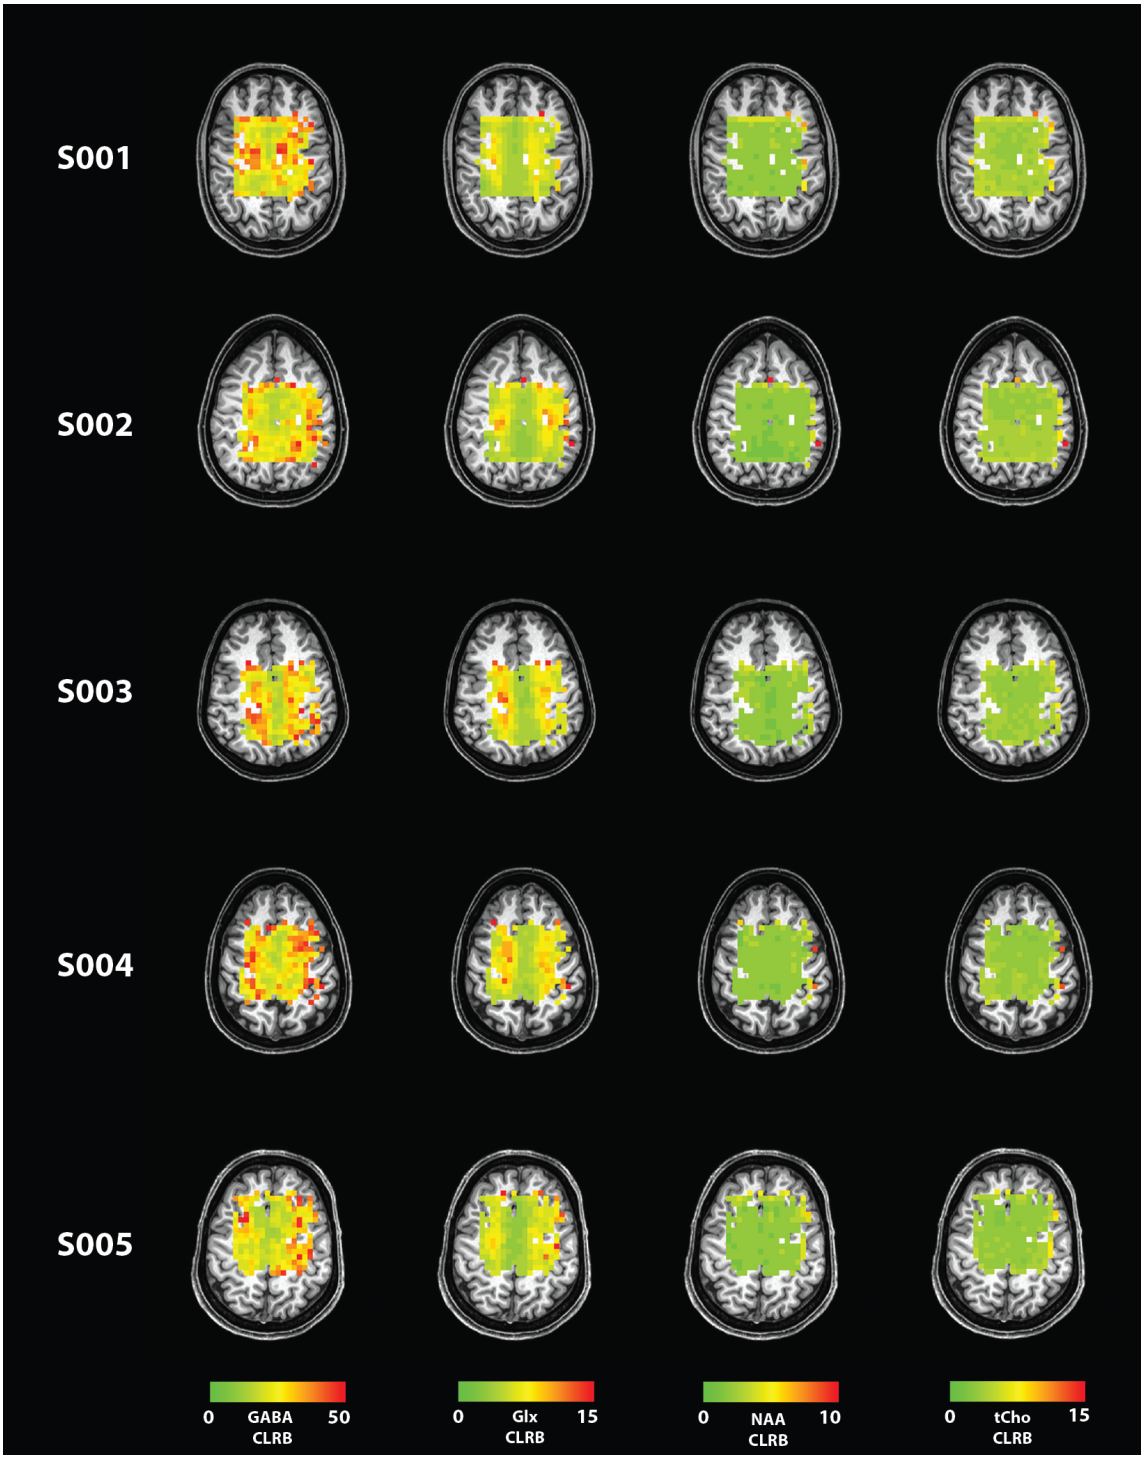


Figure 7.6. CRLB values for GABA, Glx, tNAA, and tCho corresponding to participants in Figure 7.5. As expected, CRLB values for Glx, tNAA, and tCho are consistently lower than GABA.

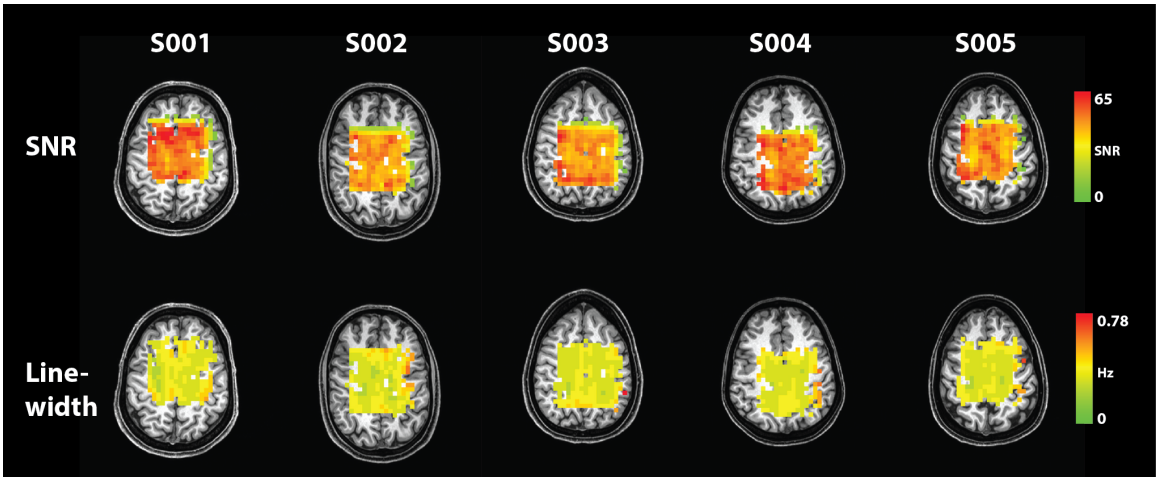


Figure 7.7. SNR and LW maps for participants shown in Figure 7.5 and Figure 7.6. SNR and LW are distributed uniformly across cortex. The band of low SNR voxels anterior and right in the slab are at the boundary of the fully-excited area and were not considered in the analysis.

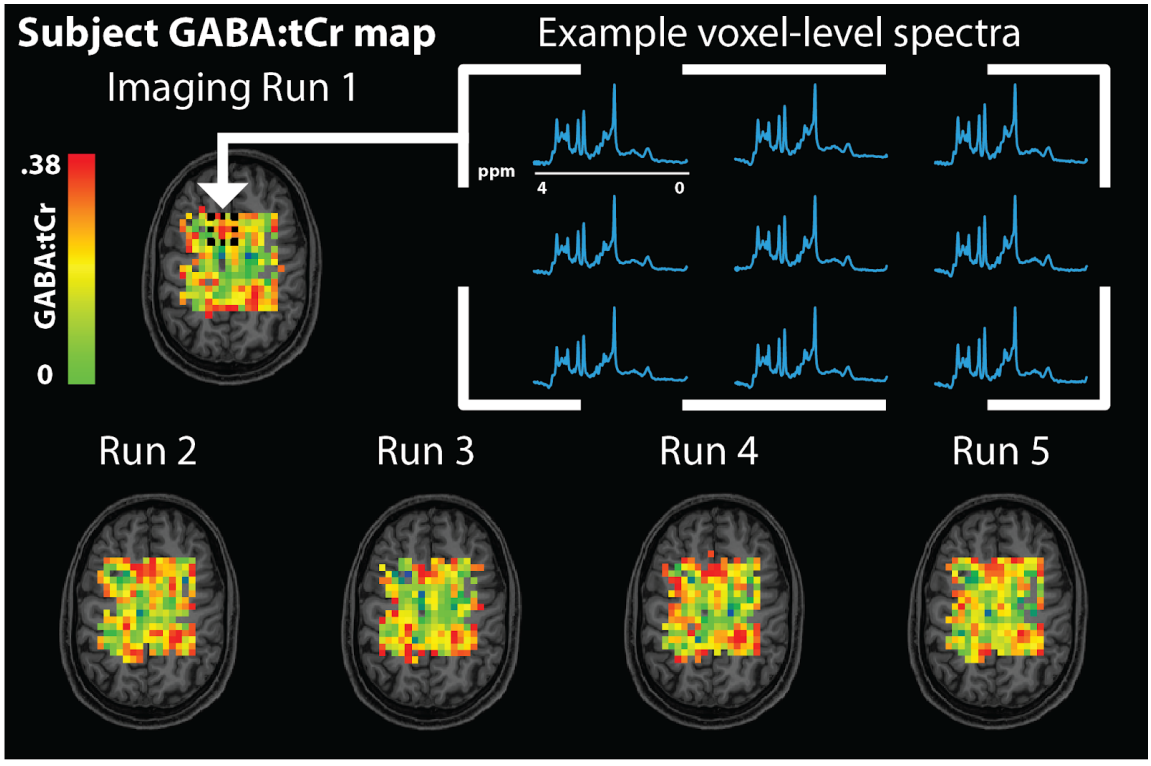


Figure 7.8. Map from each of the five representative scans for a single subject, along with representative spectra from nine contiguous, representative voxels (box, inset). GABA distribution maps were highly reproducible at the individual participant-level. The representative voxels, taken from frontal cortex, a difficult area to achieve good spectral quality, demonstrate the high-quality spectra underlying metabolite quantification.

7.4.2 MRSI reliability

One goal of the present Chapter was to establish the reliability of GABA and Glx measurements using this data acquisition and processing protocol. I quantified the reliability of the metabolite measurements in two ways. First, for all participants, I calculated similarity (Pearson's r) between the images acquired from each of the 5 imaging runs. This analysis provides a measurement of the reliability at the image-level across runs, and also provides a theoretical upper-bound on any correlation between neurochemicals and behaviour. Figure 7.9 (left) shows the distribution of correlation coefficients across imaging runs for GABA and Glx. The median and interquartile ranges of the r-values for GABA, Glx, tCho, and tNAA are shown in Table 7.1.

Second, I calculated the coefficient of variability. This analysis provides an estimate of the reliability of a measurement at the individual voxel-level and has been used previously to quantify the reliability of measurements taken using svMRS [346]. Figure 7.9 (right) shows the distribution of coefficient of variability values, and the median and interquartile ranges of the coefficient of variability values for GABA, Glx, tCho, and tNAA are shown in Table 7.1.

As expected, the distribution of r-values and coefficients of variability for GABA are lower than the other neurochemicals. However, this level of reproducibility is on par with measurements recorded using svMRS with both SPin ECho full-Intensity Acquired Localised MRS (SPECIAL) and MEGA-PRESS localisation techniques [346]. In addition, the CRLB values here are consistent with CRLB values recorded using PRESS localization at 7T [120]. This high degree of reproducibility and quality of spectral fitting provides confidence that the MRSI

measurements taken here are consistent in quality with the current gold-standard svMRS acquisitions.

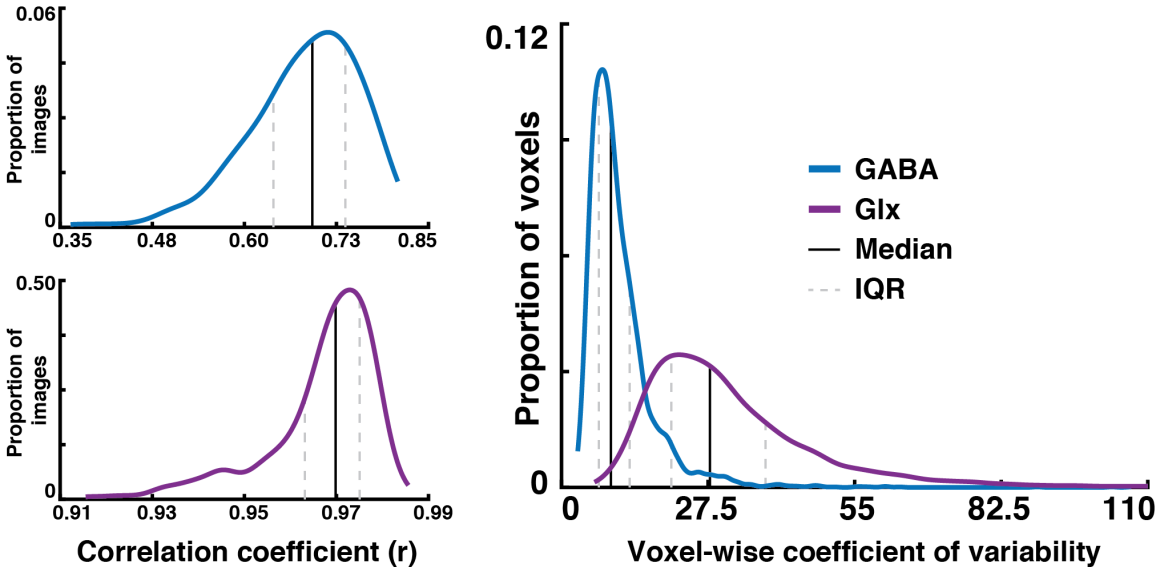


Figure 7.9. Distribution of reproducibility assessment measurements calculated in the present study. Left. For each participant, the similarity of the five GABA (blue) and Glx (purple) images to each other was calculated (number of images = 5). The average image x image correlation was higher for Glx compared to GABA. Right. The coefficient of variability for each voxel in the GABA and Glx images were calculated. In general, there were higher coefficients of variation for Glx compared to GABA. Solid black lines show median value, dotted lines show interquartile range.

Table 7.1. Quantification of data quality and reproducibility across participants.

Measurement Median (Interquartile range)	**GABA**	**Glx**	**tCho**	**tNAA**
Cramer-Rao Lower-Bound (percent)	23 (19-29)	6 (4-7)	4 (< 1)	2 (< 1)
Median voxel-wise correlation (r-value)	0.6922 (0.63-0.73)	0.969 (0.961-0.975)	0.987 (0.983-0.991)	0.992 (0.991-0.994)
Coefficient of variability	9.28 (6.96-12.77)	27.85 (20.56-38.27)	47.56 (34.95-64.45)	73.21 (55.05-99.83)

7.4.3 Characterising relationships between neurochemical concentrations and behaviour

7.4.3.1 Neurochemical distribution across parcels

My neuroscientific goal for the present study was to investigate the relationship between regional neurochemical measurements and sensorimotor performance traits in this sample. For these analyses, the 5 MRSI measurements were averaged to form the final image. In order to circumvent issues of aligning the MRSI voxels across participants, the individual participant MRSI images were warped (using nearest-neighbour interpolation) to the MNI atlas, and the average neurochemical concentration in each parcel from the Desikan atlas was calculated. Parcels were only considered if all participants had at least 50 voxels overlapping with the region of interest after alignment.

The distributions of GABA and Glx concentration from each MRSI parcel with adequate coverage in all participants are shown in Figure 7.10. Each data point represents a participant. I next examined the extent to which regional concentrations of GABA and Glx were correlated with one another (i.e., addressing whether participants with high GABA in one parcel tended to have higher GABA in other parcels). As expected, because many MRSI voxels were shared, the measurements of GABA and Glx from neighbouring parcels were significantly correlated (Figure 7.11). The parcel to parcel correlation was significantly higher for Glx than GABA (Difference between mean of z-transformed r-values for GABA versus Glx = -0.411, $t_{(54)} = -4.62$, $p < 0.001$).

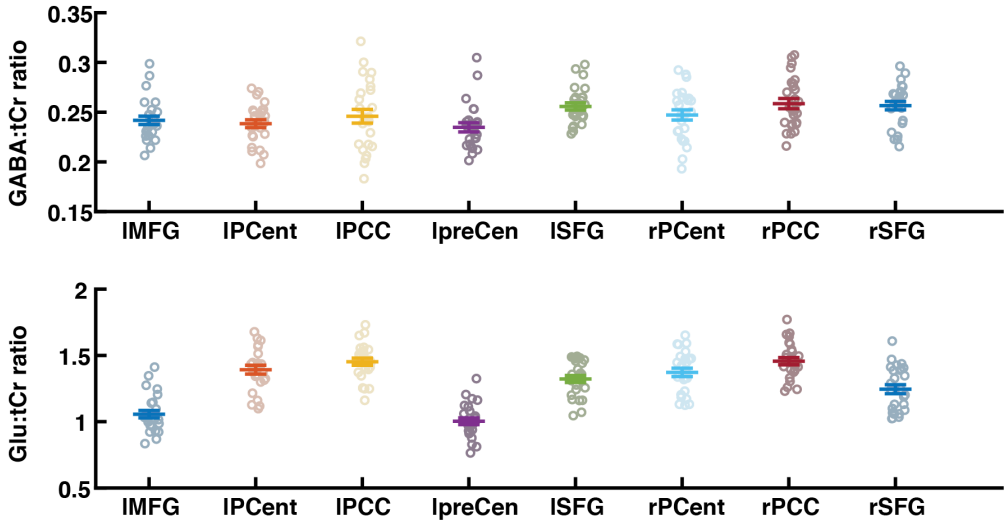


Figure 7.10 Parcel-wise GABA (upper) and Glx (lower) concentrations in each Desikan atlas parcel with adequate coverage (greater than 50 resampled MRSI voxels). Midline regions show relatively higher concentrations of GABA and Glx compared to lateral regions due to higher concentration of grey matter in MRSI voxels sampled by those parcels. lMFG: Left caudal middle frontal gyrus, lPCent: Left paracentral gyrus, lPCC: Left posterior cingulate cortex, lpreCent: Left precentral gyrus, lSFG: Left superior frontal gyrus, rPCent: Right paracentral gyrus, rPCC: Right posterior cingulate cortex, rSFG: Right superior frontal gyrus.

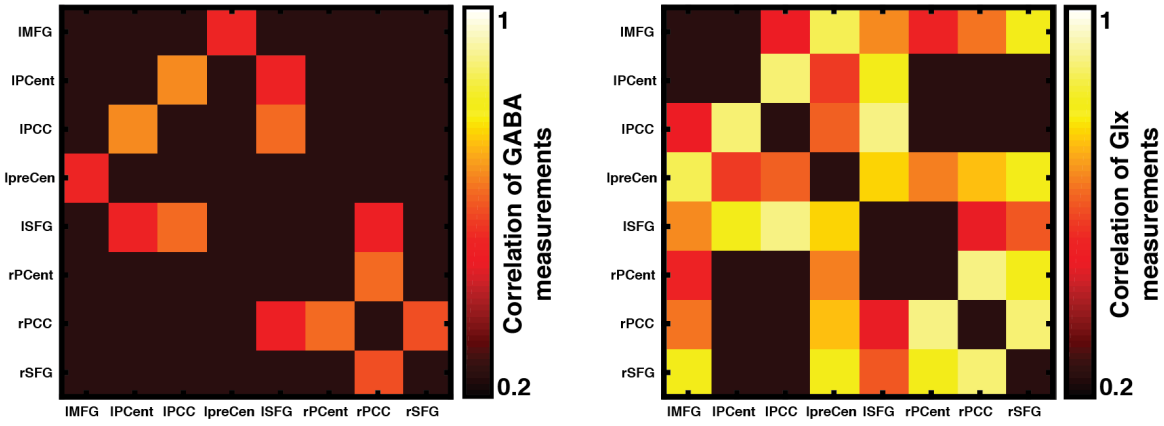


Figure 7.11. Correlation between concentrations of GABA (left) and Glx (right) across parcels. The across parcel correlation tended to be higher for Glx than GABA. MFG: Left caudal middle frontal gyrus, lPCent: Left paracentral gyrus, lPCC: Left posterior cingulate cortex, lpreCent: Left precentral gyrus, lSFG: Left superior frontal gyrus, rPCent: Right paracentral gyrus, rPCC: Right posterior cingulate cortex, rSFG: Right superior frontal gyrus.

7.4.3.2 Right superior frontal gyrus GABA concentration is associated with sensorimotor performance

Next, I sought to relate regional GABA concentration to sensorimotor performance. To do this, I used PCA summarise the participants scores on the four behavioural metrics into a single sensorimotor score [Action selection proficiency (PC weight = -0.49), Sequence learning (-0.48), Sequence chunking (0.011), and Perceptual acuity (0.72), Variance explained = 40%]. Importantly, the weights for each component suggest that a positive sensorimotor score reflects *worse* performance (i.e., action selection proficiency and sequence learning, where positive numbers indicate better performance, have a negative coefficient weights; whereas for perceptual acuity, on which lower scores indicate better performance, had a positive coefficient weight). Therefore, a *lower sensorimotor score indicates better performance*.

Having combined the behavioural tasks into a single metric, I then examined whether a participant's sensorimotor score could be related to neurochemical concentrations in any of the parcels from the Desikan atlas. To do this, I correlated (Pearson's) participants' sensorimotor scores with the corrected GABA concentrations in the parcels (corrected p-value = 0.05/8 = 0.0063). Using this approach, the only significant correlation with behaviour was with GABA concentration within the right superior frontal gyrus ($r_{(19)} = 0.67$, $p = 0.001$; Figure 7.12). This relationship suggests that lower levels of inhibition of the right superior frontal gyrus are related to better behavioural performance on these sensorimotor tasks. No other regions showed a significant relationship between GABA and behaviour, and no regions showed a significant relationship between Glx and behaviour (all p's > 0.05).

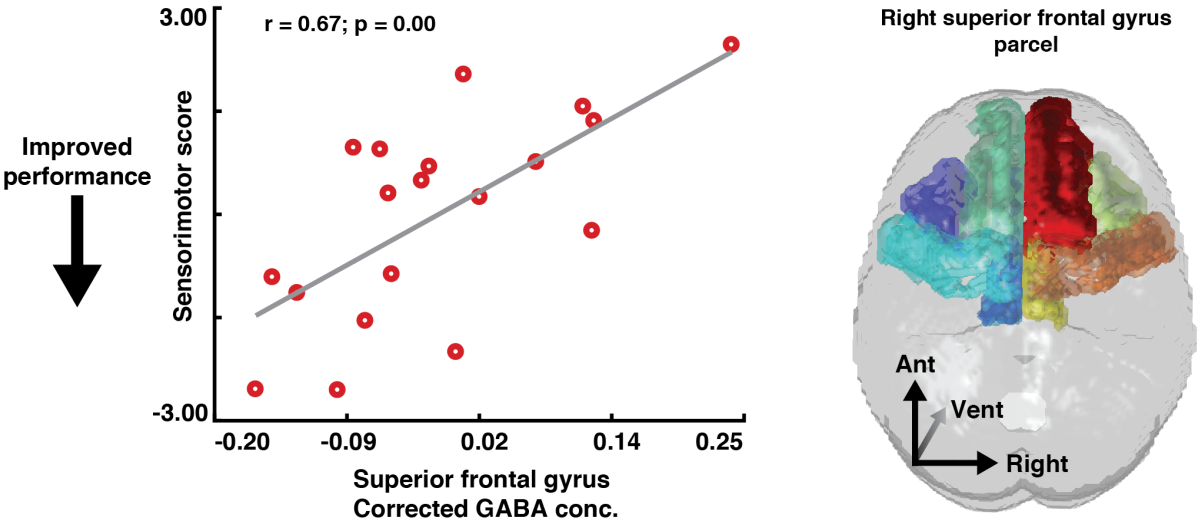


Figure 7.12. Correlation between right superior frontal gyrus (shown in red in the Right panel) GABA concentration and Sensorimotor score. Lower sensorimotor score indicates better performance.

Next, I explored the relationship between the individual behavioural metrics separately. I found no significant correlations between GABA concentration and performance when correcting for multiple comparisons across the number of regions tested. However, when considering a more liberal threshold for significance ($p < 0.05$, uncorrected; Figure 7.13), right superior frontal gyrus GABA was significantly related to both SRTT learning ($r_{(19)} = -0.544$, $p = 0.016$), and perceptual acuity ($r_{(19)} = 0.48$, $p = 0.036$). In addition, the GABA concentration in left posterior cingulate gyrus was related to perceptual acuity ($r_{(19)} = 0.48$, $p = 0.034$). There was also a strong relationship between right superior frontal gyrus GABA concentration and the RT specifically during Select trials in the action selection task such that lower GABA concentrations was related to faster RTs on Select trials ($r_{(19)} = 0.61$, $p = 0.0052$; data not shown). No correlations between Glx concentration and behavioural performance were present for any region, even at this more liberal threshold.

The relationships above were qualitatively similar when using the canonical nuisance correction strategy rather than the CRLB correction

[Correlation between right superior frontal gyrus GABA x sensorimotor performance: ($r_{(19)}$ = 0.49, p = 0.039); right superior frontal gyrus GABA x sequence learning: ($r_{(19)}$ = -0.45, p = 0.048); right superior frontal gyrus GABA x perceptual acuity: ($r_{(19)}$ = 0.29, *n.s.*); Left posterior cingulate cortex x Perceptual acuity ($r_{(19)}$ = 0.283, *n.s.*)

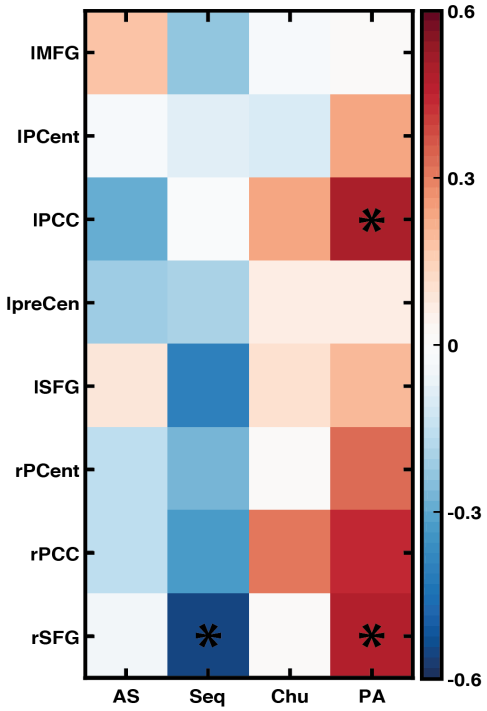


Figure 7.13 Correlation between GABA concentration in each Desikan atlas parcel with the behavioural metrics from each task. Right SFG GABA concentration was significantly correlated with SRTT learning and perceptual acuity. Asterisks (*) indicate correlations significant at p < 0.05 uncorrected. AS = Action selection proficiency, Seq = Sequence learning, Chu = Chunking, PA = Perceptual acuity. lMFG = Left middle frontal gyrus, lPCent = Left paracentral gyrus, lPCC = Left posterior cingulate cortex, lpreCen = Left precentral gyrus, lSFG = Left superior frontal gyrus, rPCent = Right precentral Gyrus, rPCC = Right posterior Cingulate Cortex, rSFG = Right superior frontal gyrus.

7.4.3.3 GABA concentration in the left attention network predicts network-level functional connectivity

One question of recent interest is the neurochemical basis for network-level functional connectivity. One potential mechanism to regulate within-

network functional connectivity strength is the inhibitory state of key nodes within a functional network. Previous studies have suggested that GABAergic inhibition in a key network node is inversely related to connectivity within the network [153, 154]. However, to date these studies have been limited in their ability to answer these questions as they have only quantified neurochemical concentrations from one voxel of interest within any given network.

To address this question, I used a group-ICA approach on the resting-state fMRI data to establish functional connectivity networks [332]. Of these, four networks had adequate MRSI coverage (defined as having at least 100 resampled MRSI voxels within the network): the left attention network, dorsal sensorimotor network, ventral sensorimotor network, and the cingulopercular network. I obtained a connectivity-strength value for each participant from each of these four networks separately using the dual-regression approach, and I looked for a relationship between network-level functional connectivity strength and GABA concentration in the network using a Pearson's correlation.

Of the networks tested, I only found a significant relationship between mean connectivity within the Left attention network (Shown in Figure 7.2) and GABA concentration. Consistent with previous studies, this relationship was negative such that higher GABA concentration was related to lower mean network-connectivity ($r_{(19)} = -0.55$, $p_{two\text{-}tailed} = 0.018$; Figure 7.14). However, this relationship did not survive Bonferroni-correction for the number of networks tested (Bonferroni corrected $p = 0.05/4 = 0.0125$). Glx concentration in the Left attention network was not related to network-level connectivity ($r_{(19)} = -0.08$, $p = 0.74$). No other networks showed a significant relationship between network-level connectivity-strength and GABA concentration (all p's > 0.17, uncorrected).

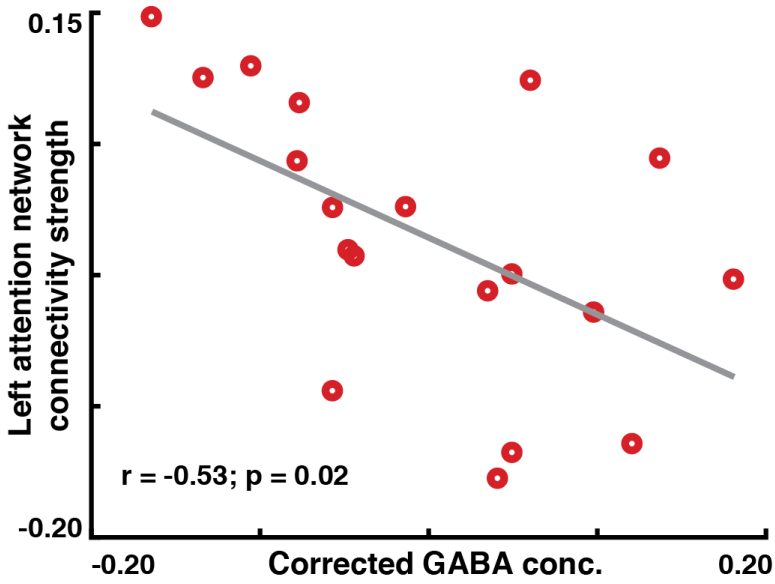


Figure 7.14. Correlation between GABA and functional connectivity strength for the left attention network. The negative relationship between GABA concentration and connectivity strength suggests that higher GABA levels leads to lower intra-network connectivity and greater inter-network connectivity of this network.

7.4.3.4 Towards a relationship between neurochemical concentration, functional connectivity and behaviour

Finally, I wanted to take these analyses a step further by linking these neural measures with behavioural performance. To this end, I related the interaction between network-level resting-state functional connectivity and neurochemical concentration to sensorimotor behaviour. Given the pattern of the results above, I fitted a linear model using both the left attention network strength and CRLB-corrected GABA concentration as predictors of sensorimotor performance. This model was significant at the level $p < 0.1$ (Full model F-statistic: $F_{(3,15)} = 2.839$, $p = 0.073$). Applying an analysis of variance to the terms of the model, only the interaction term of the model (mean connectivity x GABA concentration) significantly contributed to the model's prediction (Coefficient estimate = -0.96, $F_{(1,15)} = 6.83$, $p = 0.020$). The negative coefficient suggests that the relationship is best characterized as a U-shaped function, where a balance

between within-network GABA and connectivity strength was related to optimal behavioural performance in this sample (Figure 7.15).

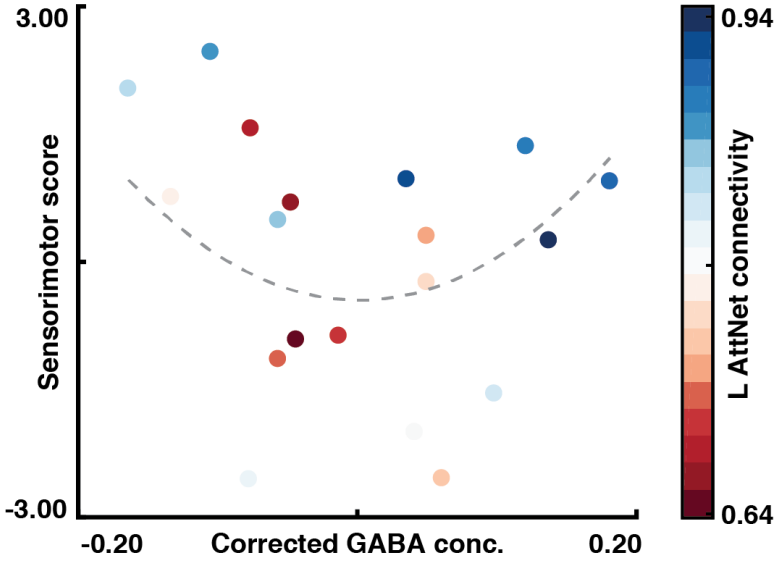


Figure 7.15. The interaction between left attention network GABA concentration and mean network-connectivity strength relates to sensorimotor score. Lower sensorimotor score indicates better participant performance. The model prediction (dotted line) shows the predicted sensorimotor score at a given GABA concentration based on the relationship between GABA and left attention network connectivity. Each point represents a single participant, coloured by the functional connectivity of the left attention network.

For completeness, I checked whether there was any correlation between left attention network GABA concentration and the GABA concentration of the Desikan atlas parcels, in particular the left and right superior frontal gyri. There was no significant relationship between the GABA concentrations in the left attention network GABA and right superior frontal gyrus GABA ($r_{(19)}$ = -0.16, p = 0.5) and only a weak relationship between the GABA measurements in the left attention network and left superior frontal gyrus ($r_{(19)}$ = -0.412, 0.079). Among the Desikan atlas parcels with adequate coverage, GABA concentration in the left

middle frontal gyrus had the strongest relationship to the left attention network GABA concentration ($r_{(19)} = 0.568$, $p = 0.011$).

7.5 Discussion

In the present study, I used an advanced MRSI approach to quantify neurochemicals in order to relate GABA concentration to network-level functional connectivity, and further, to sensorimotor behaviour. I first demonstrated that this novel MRSI technique yields reproducible GABA and Glx measurements, and then developed an analysis pipeline for investigating the relationship between regional neurochemical concentrations and behaviour. The application of this MRSI acquisition strategy to address a neuroscientific question marks a distinct technical advance and opens up new possibilities for studying network processes from a neurochemical perspective.

7.5.1.1 A potential link between neurochemical concentration, functional connectivity, and behaviour

The main result of this study is a potential link between the left attention network GABA concentration, functional connectivity strength, and sensorimotor skill. I found that GABA concentration of the left attention network was negatively correlated with functional connectivity within that network, consistent with prior literature showing that found higher GABA concentrations in a resting-state network predicted lower within network-level functional connectivity in the motor network [153] and default mode network [154]. Additionally, the relationship between GABA and resting-state functional connectivity of the left attention network predicted participants' sensorimotor behaviour, such that participants with high GABA performed better if the they had lower within-

network connectivity, and participants with low GABA performed better if the they had higher within-network connectivity. Others have speculated that MRS-measured GABA concentration reflects local inhibitory processes [347]. This leads to the hypothesis that a predisposition for balanced local-inhibitory tone and within-network connectivity strength leads to better sensorimotor performance. Future work may test this hypothesis using pharmacological manipulations or non-invasive brain stimulation to modulate the local inhibitory state and network-level functional connectivity.

7.5.1.2 Right superior frontal gyrus GABA concentration is associated with sensorimotor performance

In addition to testing network-level GABA concentration, I also examined whether concentration of GABA and Glx within anatomically-defined regions of interest was related to sensorimotor performance. I found that GABA levels in the right superior frontal gyrus were strongly related to overall sensorimotor ability. The superior frontal gyrus region of interest in the Desikan atlas is quite large, overlapping with preSMA and the anterior dorsal PMC, which are critical regions for high-level sensorimotor behaviours such as sequencing [9] and action selection [52, 344]. Although the tasks were performed with the right hand, right PMC is known to play a dominant role in sequencing and action selection [55, 344], and therefore it is not surprising that the neurochemical concentration of this region was related to sensorimotor performance.

Others have found that sensorimotor cortex GABA concentration relates to sensorimotor behaviours such as reaction time [348] and perceptual acuity [225], but I did not observe this relationship. This may be due to the coverage of the

imaged region in this study: while I did have some coverage of precentral gyrus, the "hand knob" was not covered in the majority of participants. This choice was necessitated by skull geometry, as positioning the acquisition area for the MRSI sufficiently laterally to cover the precentral gyrus precluded covering other areas of interest, such as medial and lateral prefrontal cortex. Based on data obtained in this experiment, the placement of the imaged area can be optimized in the future.

7.5.1.3 Questions raised by the overall pattern of results

The overall pattern of results raise several important questions that might be addressed in future studies. The first question concerns the ability to relate neurochemical concentrations to resting-state network-level connectivity, and why I only found a relationship between these variables in the left attention network. Several factors in the study design might address this question.

One factor is the coverage of the resting-state networks by the area acquired for MRSI. Theoretical papers suggest that network-level connectivity strength may be driven by some nodes in the network more than others [202]. It is therefore possible that the acquired area for MRSI did adequately cover these driver nodes such that the neurochemical measurements could predict connectivity-strength in all networks considered in this study. To this point, it is important to note that the two other studies examining the relationship between neurochemical concentration and network-level connectivity used svMRS, and those researchers were able to place the MRS voxel in areas of the network of interest that presumably drive network activity [153, 154].

Another possible factor leading to a relationship between the left attention network GABA and connectivity-strength is the vigilance task participants

performed during MRSI acquisition. Because the acquisition duration was 45 minutes, I asked participants to perform a vigilance task that required pressing a button with the right index finger. This task may have driven activity in the left attention network, as participants prepared themselves to make responses, which may have modulated GABA levels in this network. It has previously been found that many tasks modulate GABA levels in task-relevant areas, for example in the motor cortex during motor learning [16, 120], lateral prefrontal cortex during memory encoding [349], and anterior cingulate cortex during statistical learning [350]. In addition, the connectivity of the attention network, which includes dorsolateral prefrontal cortex, middle cingulate cortex, and posterior parietal cortex, has been found to predict performance on attention tasks [351, 352]. Thus, implementing a vigilance task may have sensitized this experiment to find this specific relationship.

The second question raised by the overall pattern of the results is the laterality difference in the relationship to behaviour found between the two analyses. For the anatomical atlas, I found a relationship between the right superior frontal gyrus GABA concentration and sensorimotor proficiency. On the other hand, for the resting-state defined networks, I found that the interaction between left attention network GABA concentration and network-level connectivity strength related to behavioural performance. The discrepancy between these tests is likely due to the resting-state network GABA considering voxels in lateral prefrontal cortex that were not located in the superior frontal gyrus. Consistent with this explanation, there was no significant correlation between right superior frontal gyrus GABA concentration and left attention network GABA concentration, and only a weak negative correlation between

GABA concentration in the left attention network and GABA concentration in left superior frontal gyrus. Thus, these two approaches are sensitive to the neurochemical concentrations in different regions and likely reflect different neural processes.

7.5.1.4 Semi-LASER MRSI produces reproducible GABA measurements

In addition to the neuroscientific result, the second major advance I present in this chapter is methodological; demonstrating that the non-water suppressed semi-LASER with DW-CRT acquisition can produce high-resolution maps of GABA and Glx distribution in the human brain. While I have previously shown that we could use this technique in a clinical setting, in that demonstration I did not try to quantify GABA concentration [204]. Here, I adjusted the acquisition strategy to make quantification of GABA possible and quantified the reproducibility of our measurements over five image acquisitions in 26 participants. We achieved reproducibility comparable to the current gold-standard unedited svMRS technique at 3T for GABA detection [346], but with approximately 41 times improvement in resolution (0.375 mL in our study compared to 15.625 mL in [346]). In the study by Near and colleagues, the average coefficient of variability for svMRS GABA was 8.7%, and in our study, the reproducibility for GABA was 9.28%. This level of reproducibility is encouraging and suggests that future studies could acquire fewer MRSI repetitions, freeing up valuable time for research scanning.

The final major accomplishment of the present study was the implementation of a pipeline for analysis of MRSI data. One key issue for MRSI with a single slice is the alignment across participants. One option for group level

statistics is to do voxel-wise statistics, with all participants registered to standard space, as is common for other types of MRI data. However, because there is no variation in the z-axis, it is not clear whether a given voxel in standard space would be representative for each participant. Instead, I aligned participants to a standard atlas and averaged across anatomically defined parcels (Desikan atlas). This approach allowed me to determine a threshold for adequate representation in all participants for each parcel, and to account for individual differences in anatomy as well as partial voluming across these regions. In addition, by using a standard atlas, I avoid any biases that might arise from manual selection of ROIs, and makes the ROIs used here easily relatable to other studies.

7.5.2 Limitations

By applying MRSI in an experimental setting, the present study constitutes a significant technique advance. However, there are several limitations that should be noted. One notable limitation is the relatively small number of participants for a correlational study. All correlations are thus interpreted with this in mind. Nevertheless, these results help generate hypotheses for future studies. Second, although acquisition of MRSI constitutes a massive advance in both MRS resolution and number of voxels acquired, some regions of the brain were still not densely sampled. Development of simultaneous multi-slice acquisitions [353-355] will be another technical advance to consider for investigations in the future.

7.5.3 Open questions and future directions

One interesting feature of the data that has not been fully explored is the covariation of neurochemicals across different anatomical regions. Anatomical covariation, for example cortical thickness and T1:T2 ratio is an important

phenotype that varies across individuals and may be a marker of brain function in health and disease [356-358]. Understanding neurochemical covariation may be another important dimension on which to stratify populations.

In addition, the present study only considered univariate relationships between a single neurochemical and a single behaviour or a single resting-state network. One factor not considered was how the covariation among regional neurochemical concentrations might relate to the covariation in brain activity across multiple regions or in multiple networks. These multivariate relationships have been used successfully to link patterns of brain function to behavioural traits [359], and I intend to investigate this as an approach in future analyses.

Finally, although I related variation in neurochemical concentration to the *ability to learn* demonstrated in a separate session, I did not measure neurochemical variation *during learning*. An important next step will be to understand how inhibitory processes vary across the brain during learning.

7.6 Conclusions

In summary, in the present study I made progress in linking regional GABA concentration to functional connectivity and sensorimotor performance. In addition, I demonstrated a novel MRSI imaging technique produces reliable metabolite maps and show that this technique is suitable for use in a basic science setting. This study paves the way for further investigations using this tool to investigate the molecular mechanisms that may underlie cortical plasticity at the network-level in humans.

8 General discussion

In this chapter, I will outline the major findings presented in this thesis. First, I summarize the key findings and main conclusions from each chapter. Second, I briefly revisit the limitations of each study, and discuss these limitations in context of plasticity investigations. Finally, I discuss possible directions for future research and highlight general conclusions that can be drawn from this thesis.

8.1 Summary of findings

In this thesis, I investigated the behavioural and neural impact of motor skill learning using a combination of longitudinal fMRI, behavioural manipulations, and MRSI.

In Chapter 3, I investigated the extent to which changes in functional connectivity after training were specific to task. I used a longitudinal design that allowed me to directly compare the impact of visuospatial and motor sequence learning on functional connectivity. I found that the changes induced by training were remarkably specific to the trained task: hippocampal functional connectivity was primarily modulated by visuospatial training, whereas sensorimotor cortex functional connectivity was primarily modulated by training on the motor sequencing task. In addition, I observed non-specific changes in functional connectivity that occurred even in the absence of task performance. Overall, this study demonstrates that functional connectivity is sensitive to task specific changes and highlights the importance of adequate controls in the context of plasticity research.

In Chapters 4 and 5, I investigated the impact of reward and punishment on behavioural performance during skill learning. In Chapter 4, I first focused on the extent to which task demands impact the effect of reward and punishment on performance. I trained participants on two different skills, the FTT and the SRTT, augmented with valenced feedback. I found that the impact of reward and punishment depended on task demands: punishment improved performance on the SRTT, but impaired performance on the FTT. In addition, I did not observe any effect of reward or punishment on memory formation. Thus, in this study, I showed that the simple heuristic widely applied in motor skill learning, that punishment may benefit performance while reward may benefit memory formation, and therefore any benefit of reward to memory formation (e.g., in adaptation [151, 266]) or associative learning (e.g., [360]) may specific to those tasks.

In Chapter 5, I then investigated whether the intention of the participant to learn interacts with the effect of valenced feedback, focusing on the SRTT. I found that intentionality modulated the impact of feedback on performance during learning. When participants were learning unintentionally, punishment quickened sequence knowledge expression, but not when participants intentionally tried to learn the sequence. The results of this study suggest that the effect of feedback may interact with higher cognitive processes, such as intention, which further complicates any attempt to establish general principles that explain the effect of reward and punishment on skill performance.

I next examined the effect of reward and punishment on brain activity after training using resting-state fMRI in Chapter 6. Based on the behavioural data presented in Chapter 4, I anticipated that reward and punishment would

differentially impact brain activity changes due to training, but that these effects would depend on the task performed. Indeed, I found that reward and punishment had dissociable effects on resting-state functional connectivity and confirmed that the regions implicated depended on the task being learned. These results are consistent with the hypothesis that the effect of feedback is dependent on the task performed, perhaps due to the neural circuitry involved in task performance.

Finally, in Chapter 7, I undertook a preliminary investigation into the relationship between network-level neurochemical distribution, functional connectivity, and motor performance traits. For this study, I applied a novel MRSI technique to quantify the distribution of neurochemicals across the brain. I showed that this technique gives highly reproducible measurements of the neurotransmitter GABA, and that regional GABA concentrations were related to sensorimotor performance. Further, I provided preliminary evidence for a link between network-level GABA concentrations, network-level functional connectivity, and sensorimotor performance. This study provides the foundation that future studies may use MRSI to investigate the large-scale neurochemical consequences of learning.

8.1.1 Overall summary

In summary, my aims in this thesis were to elucidate the consequences of short-term motor skill training on human brain activity, and to move toward a physiological account of the functional plasticity. These aims were primarily motivated by animal models, which suggest that activity changes within brain networks can produce robust short-term behavioural changes, and that the activity changes occur broadly across the brain [113, 140]. Indeed, this is what I found; functional connectivity was broadly impacted after task performance. In

addition, the observed effects were highly task-specific. These data suggest that the network-level changes induced by training do not trivially generalize across tasks. These data underscore the importance of considering results in light of the task being performed, and the importance of resisting the temptation to make general claims when only a single task context is considered.

8.2 General limitations

The specific limitations of each study were discussed in the corresponding chapters. In this section, I discuss general limitations of the approaches I took to examine functional plasticity.

8.2.1 Limitations of resting-state functional connectivity

8.2.1.1 Interpreting resting-state functional connectivity changes

As described in Chapter 2, resting-state functional connectivity has been widely used to study brain organization. However, the biological basis of functional connectivity is still not well-understood. Thus, although I observed significant functional connectivity changes after learning in the studies presented in Chapters 3 and 6, the underlying physiological changes that took place are not known. Animal models, detailed in Chapter 2, provide compelling evidence that the spontaneous BOLD-signal fluctuations reflect, at least in part, ongoing neural activity. Thus, it is tempting to conclude that changes in functional connectivity between brain regions may reflect changes in neuronal communication between these regions.

More direct investigations into the relationship between functional connectivity and ongoing neuronal activity must be undertaken in order to draw firm mechanistic conclusions about functional plasticity from fMRI-derived

functional connectivity data. Studies such as [178, 361] have provided an important first step toward understanding the relationship between ongoing neural activity and the haemodynamic response. However, functional plasticity studies critically depend on understanding how *changes* in the ongoing neural activity between two regions would be reflected in the haemodynamic responses in these regions. To this end, performing a study whereby neural activity and cerebral blood flow changes are simultaneous measured before and after learning would be a highly informative endeavour. Extending the work presented in Chapter 7 to consider the neurochemical changes during learning could also help address this limitation.

Thus, based on the studies presented in this thesis, I cannot directly implicate neuronal communication as the cause of the functional plasticity I observed, nor can I make mechanistic inferences about specific patterns of activity (e.g., hippocampal theta-oscillations) underlying the changes in BOLD-signal co-fluctuation between regions. Despite these limitations, the data presented in Chapters 3 and 6 provide important information about the neural basis of memory formation. These studies show definitively that using resting-state functional connectivity change to measure functional plasticity can provide information about the specific context in which certain anatomical regions might be involved (i.e., what contexts, such as tasks or features, engage particular brain areas). The precise roles of the implicated brain areas can then be queried with further experiments.

8.2.1.2 Correlation between functional connectivity changes and behavioural changes

While functional connectivity can provide important contextual and spatial information about the brain areas involved in memory processes, inferring a causal relationship between the functional connectivity change measured with fMRI and the behavioural improvements is challenging. Showing specificity to the trained task (Chapter 3) or to condition (Chapter 6) helps make any observed changes more compelling but does not resolve this issue entirely.

Demonstrating a relationship between any functional connectivity change after training and behavioural performance may provide evidence for a causal relationship. I was not able to demonstrate any relationship between functional connectivity change and behavioural performance in these studies. Therefore, it may be that the changes in functional connectivity I observed were related to *task performance* (e.g., rumination on the task [259]) but may have been epiphenomenal to *memory formation.* The lack of a behavioural relationship does not prove that the observed changes are not relevant to memory formation. For example, it may be that functional connectivity changes and behavioural changes are not linearly related, or that the behavioural metric being used is a poor measure of what was 'learned.'

In addition, in Chapter 3, I observed several regions that showed a correlation between resting-state functional connectivity change and behavioural change but *no task-specific change in rsFC*. This relationship is also difficult to interpret. For example, it could be that the region showing this effect acts as a switch: *if this change occurs, then a participant will learn*. However, another possibility is that these relationships are not robust and would not replicate.

This is particularly problematic in the studies presented in this thesis due to the relatively low sample size (Chapter 3), particularly for correlational studies. Indeed, I did not investigate relationships between functional connectivity change and behaviour in the study presented in Chapter 6 because I could not be confident in any result discovered due to the low sample size within each group.

Thus, I cannot conclude that the functional connectivity changes observed in this thesis were causally related to memory formation. However, these studies may provide a basis for testing the casual role of a region in these contexts using interventional techniques, such as transcranial brain stimulation techniques.

8.2.1.3 Generalizability of findings

Task is rarely included as a factor in fMRI studies, and yet, the generalizability of findings to other tasks or to external circumstances is often implied. One important aspect of this thesis was demonstrating that, for both behavioural and neural effects of training, many changes were task-specific. However, while using two (Chapters 3-6) or three (Chapter 7) tasks did allow me to address the question of generalizability of my results, this was by no means a comprehensive investigation into these effects. As such, it would be inconsistent with these data to suggest that my results, such as the benefit of punishment to early skill memory formation, would generalise beyond the scope to the presented work.

Failure to generalize to new contexts is a major issue in psychology and cognitive neuroscience [362]. Because I incorporated multiple tasks, the results I presented could highlight this important issue in both behaviour and imaging experiments. This demonstration may help motivate others to consider

generalizability tests in future studies and to promote the use of experimental paradigms with superior generalizability.

8.2.2 Limitations of behavioural tasks

8.2.2.1 Response time as an outcome measure

In addition, several of the behavioural tasks I employed (Chapters 3-5 and 7: SRTT, Chapter 7: AST) use RT as an outcome measure. RT is highly variable across individuals but also within an individual in different contexts [363-365]. This concern is partially mitigated by establishing a baseline measure in all of the studies, but the variability of reaction time can further complicate the generalizability of findings, discussed above. In addition, RTs are not subserved by a single physiological process, but rather reflect several processes, like preparatory state, computation speed, cortical excitability, etc. Thus, the use of RTs as an outcome measure cannot disambiguate performance and knowledge. Accordingly, I observed that participants who learned with reward and punishment evidenced the same level of memory formation despite having differences in RT during training (Chapter 4 and 5).

RT is a common outcome measure used in many experiments. Therefore, using RT was a compromise to make my studies more relatable to prior work in the field. In addition, the majority of these tasks were conducted in the scanning environment, which limited the breadth of tasks that might be performed, particularly in motor learning. Thus, choosing RT as an outcome measure was an important compromise. Future studies might make use of the separable components of RT [363]. For example, decomposing RT into the execution and

preparation components may be a fruitful next step towards determining a mechanistic explanation as to the benefit of punishment to SRTT performance.

8.3 Future directions and open questions

In this thesis, I have investigated both behavioural and neural consequences of short-term training. These studies, while limited in scope, have raised many important questions that can serve as the impetus for future work. A few of these possibilities are presented below. I begin by discussing my plans for extending the findings presented in Chapter 7, before turning to possible future studies.

8.3.1 How does neurochemical distribution across multiple regions predict task performance?

In Chapter 7, I presented preliminary evidence for a link between neurochemical concentration and network-level functional connectivity. As part of that study, I demonstrated the ability to obtain reproducible GABA measurements across many brain regions. However, I did not leverage this technique to explore how the neurochemical distribution across brain regions related to behaviour beyond univariate relationships. This leaves open the possibility that the covariation of neurochemicals across multiple regions could be an important predictor of behavioural performance. Similarly, the relationship between neurochemical distribution across regions and functional connectivity could be explored in more detail.

8.3.2 How does learning impact neurochemical concentration across brain?

As described in Chapter 1, within a population of neurons, the initial phase of motor skill learning is characterized by an increase in the variability of the neurons that activate during movement execution. This increase in neural variability occurs at the same time as the GABA concentration motor cortex decreases, which has been observed using MRS [16, 120]. However, as I suggested in Chapter 1, the increase in population response variability occurs in the context of a complex network of interacting brain regions with distinct neural populations. Whether the decrease in GABA, which might serve as a marker for population-level exploration, is regionally specific or whether this occurs across the whole-brain during skill learning remains unknown.

The MRSI method described in Chapter 7 is well-suited to address this question. Given the acquisition duration for the MRSI volume, an experiment could be implemented in a manner consistent with both Floyer-Lea and colleagues and Kolasinski and colleagues [16, 120]. In addition, unlike the study presented in Chapter 7, the placement of the imaged area could be optimized to target aspects of the motor network involved in task performance, including dorsal PMC, SMA, and primary motor cortex. By expanding the view of motor learning beyond primary motor cortex, a study like the one proposed would address critical limitations of prior work and greatly expand our knowledge of network-level plasticity.

8.3.3 Strategies employed when learning with reward and punishment

One important conclusion from Chapters 4 through 6 was that reward and punishment do, generally, elicit opposing behavioural and neural effects during learning, and further investigation could address the mechanism driving these effects. For example, in the experiments presented in Chapters 4 and 5, I observed a benefit of punishment to unintentional SRTT learning. One explanation for this benefit is the recruitment of attention-dependent mechanisms during performance. By integrating eye-tracking, the role of attention and implicit learning could be explored in greater depth. Anticipatory eye-movements may be an indicator about attentional processes, and I would hypothesize that participants training with punishment should be more prone to making such eye movements early in during training. In addition, pupil dilation tracks the concentration of norepinephrine and acetylcholine levels and may provide further insight into the attentional processes involved in learning with punishment [366, 367].

Another question raised by the studies in Chapters 4 through 6 was the role of dopamine when learning with reward and punishment. This question could be addressed in several ways. Blink rate may be a good proxy for dopaminergic activity [368]. Implementing a study like the ones in Chapters 4 and 5 with eye-tracking would allow measurement of this physiological signal and could useful for investigating the role of dopamine concurrent with the investigation into attentional processes presented above. In addition to eye-tracking, the role of dopamine in learning with reward and punishment could also be investigated with pharmacological manipulations. There is evidence that dopamine does not

play a role in feedback processing during adaptation learning [369], but, as discussed in Chapter 4, adaptation learning and SRTT learning are distinct processes. Because of the discrete trial structure and statistical nature of SRTT learning, altering dopaminergic activity may be more likely to impact this task. Consistent with this hypothesis, patients with Parkinson's disease may have deficits in SRTT learning [370] but not adaptation learning [371], suggesting a greater role for dopamine in SRTT learning. Thus, the role of dopamine in SRTT learning, and how this is modulated by reward and punishment, warrants further investigation.

8.3.4 Naturalistic investigations in motor learning

Collectively Chapters 4 through 6 highlighted the importance of generalizability of results to new contexts. As described above, generalizability is a critical limitation for many psychology and neuroscientific studies [362]. One aspect that limits the generalizability of results is the tight control used in experimental research. By using more naturalistic experimental designs, generalizability can be improved. While naturalistic investigations of motor learning have been difficult because of technical limitations, new advances in portable motion/eye-tracking technologies [372] and motion tracking from video recordings [373] has made such investigations more approachable.

One of the most interesting things that is possible with naturalistic investigations that is not well explored is how experts can make slight adjustments to their skills. For example, despite having thrown American footballs their entire lives, often when transitioning into the National Football League, quarterbacks will be required to make subtle adjustments to their throwing motion to better suit the professional-level (e.g., Christian Hackenberg, New York

Jets 2018 [374]). This degree of expertise is not easily recreated in the laboratory setting, and therefore we know very little about how learning interacts with such a high-degree of expertise. Investigating skill learning in 'true' experts in a naturalistic setting will certainly provide a richer understanding into the mechanisms of this process and may help to inform our real-world teaching practices.

References

1. Nabavinik, M., et al., *Especial Skills in Experienced Archers.* Journal of Motor Behavior, 2018. **50**(3): p. 249-253.
2. MacKay, D.G., *The problems of flexibility, fluency, and speed–accuracy trade-off in skilled behavior.* Psychological Review, 1982. **89**(5): p. 483.
3. Reis, J., et al., *Noninvasive cortical stimulation enhances motor skill acquisition over multiple days through an effect on consolidation.* Proceedings of the National Academy of Sciences of the United States of America, 2009. **106**(5): p. 1590-1595.
4. Krakauer, J.W. and P. Mazzoni, *Human sensorimotor learning: adaptation, skill, and beyond.* Current opinion in neurobiology, 2011. **21**(4): p. 636-644.
5. Karni, A. and D. Sagi, *The time course of learning a visual skill.* Nature, 1993. **365**(6443): p. 250-2.
6. Moore, D.R., S. Amitay, and D.J. Hawkey, *Auditory perceptual learning.* Learn Mem, 2003. **10**(2): p. 83-5.
7. Karmarkar, U.R. and D.V. Buonomano, *Temporal specificity of perceptual learning in an auditory discrimination task.* Learn Mem, 2003. **10**(2): p. 141-7.
8. Karni, A., et al., *The acquisition of skilled motor performance: fast and slow experience-driven changes in primary motor cortex.* Proceedings of the National Academy of Sciences of the United States of America, 1998. **95**(3): p. 861-868.
9. Yasuda, M., S. Yamamoto, and O. Hikosaka, *Robust representation of stable object values in the oculomotor Basal Ganglia.* J Neurosci, 2012. **32**(47): p. 16917-32.
10. Heitz, R.P., *The speed-accuracy tradeoff: history, physiology, methodology, and behavior.* Frontiers in Neuroscience, 2014. **8**(150).
11. WOODWORTH, R.S., *THE ACCURACY OF VOLUNTARY MOVEMENT.* The Journal of Nervous and Mental Disease, 1899. **26**(12): p. 743-752.
12. Bastian, A.J., *Learning to predict the future: the cerebellum adapts feedforward movement control.* Current Opinion in Neurobiology, 2006. **16**(6): p. 645-649.
13. Himann, J.E., et al., *Age-related changes in speed of walking.* Med Sci Sports Exerc, 1988. **20**(2): p. 161-6.
14. Ko, S.U., J.M. Hausdorff, and L. Ferrucci, *Age-associated differences in the gait pattern changes of older adults during fast-speed and fatigue conditions: results from the Baltimore longitudinal study of ageing.* Age Ageing, 2010. **39**(6): p. 688-94.
15. Floyer-Lea, A. and P.M. Matthews, *Distinguishable brain activation networks for short- and long-term motor skill learning.* J Neurophysiol, 2005. **94**(1): p. 512-8.
16. Floyer-Lea, A., et al., *Rapid modulation of GABA concentration in human sensorimotor cortex during motor learning.* J Neurophysiol, 2006. **95**(3): p. 1639-44.
17. Abe, M., et al., *Reward improves long-term retention of a motor memory through induction of offline memory gains.* Curr Biol, 2011. **21**(7): p. 557-62.
18. Steel, A., et al., *The impact of reward and punishment on skill learning depends on task demands.* Sci Rep, 2016. **6**: p. 36056.
19. Nudo, R.J. and R.B. Masterton, *Descending pathways to the spinal cord, III: Sites of origin of the corticospinal tract.* J Comp Neurol, 1990. **296**(4): p. 559-83.
20. Penfield, W. and E. Boldrey, *Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation.* Brain, 1937. **60**(4): p. 389-443.
21. Sanes, J.N. and J.P. Donoghue, *Plasticity and primary motor cortex.* Annual Review of Neuroscience, 2000. **23**: p. 393-415.
22. Graziano, M.S.A., C.S.R. Taylor, and T. Moore, *Complex movements evoked by microstimulation of precentral cortex.* Neuron, 2002. **34**(5): p. 841-851.

23. Graziano, M.S.A. and T.N. Aflalo, *Mapping behavioral repertoire onto the cortex.* Neuron, 2007. **56**(2): p. 239-251.
24. Kwan, H.C., et al., *Spatial organization of precentral cortex in awake primates. II. Motor outputs.* Journal of Neurophysiology, 1978. **41**(5): p. 1120-1131.
25. Lotze, M., et al., *fMRI evaluation of somatotopic representation in human primary motor cortex.* Neuroimage, 2000. **11**(5 Pt 1): p. 473-481.
26. Sanes, J.N., et al., *Shared neural substrates controlling hand movements in human motor cortex.* Science, 1995. **268**(5218): p. 1775-1777.
27. Ejaz, N., M. Hamada, and J. Diedrichsen, *Hand use predicts the structure of representations in sensorimotor cortex.* Nature Neuroscience, 2015. **18**(7): p. 1034-1040.
28. Sessle, B.J. and M. Wiesendanger, *Structural and functional definition of the motor cortex in the monkey (Macaca fascicularis).* J Physiol, 1982. **323**: p. 245-65.
29. Donoghue, J.P., S. Leibovic, and J.N. Sanes, *Organization of the forelimb area in squirrel monkey motor cortex: representation of digit, wrist, and elbow muscles.* Exp Brain Res, 1992. **89**(1): p. 1-19.
30. Nudo, R.J., et al., *Neurophysiological correlates of hand preference in primary motor cortex of adult squirrel monkeys.* J Neurosci, 1992. **12**(8): p. 2918-47.
31. Leo, A., et al., *A synergy-based hand control is encoded in human motor cortical areas.* eLife, 2016. **5**.
32. Nudo, R.J., et al., *Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys.* The Journal of Neuroscience, 1996. **16**(2): p. 785-807.
33. Peters, A.J., S.X. Chen, and T. Komiyama, *Emergence of reproducible spatiotemporal activity during motor learning.* Nature, 2014. **510**(7504): p. 263-267.
34. Wise, S., *The primate premotor cortex: past, present, and preparatory.* Annual Review of Neuroscience, 1985. **8**(1): p. 1-19.
35. Kurata, K. and E. Hoshi, *Movement-related neuronal activity reflecting the transformation of coordinates in the ventral premotor cortex of monkeys.* Journal of Neurophysiology, 2002. **88**(6): p. 3118-3132.
36. Kurata, K. and J. Tanji, *Premotor cortex neurons in macaques: activity before distal and proximal forelimb movements.* The Journal of Neuroscience, 1986. **6**(2): p. 403-411.
37. Rizzolatti, G., M. Matelli, and G. Pavesi, *Deficits in attention and movement following the removal of postarcuate (area 6) and prearcuate (area 8) cortex in macaque monkeys.* Brain: A Journal of Neurology, 1983. **106 (Pt 3)**: p. 655-673.
38. Luppino, G., et al., *Largely segregated parietofrontal connections linking rostral intraparietal cortex (areas AIP and VIP) and the ventral premotor cortex (areas F5 and F4).* Experimental Brain Research, 1999. **128**(1-2): p. 181-187.
39. Duhamel, J.R., et al., *Spatial invariance of visual receptive fields in parietal cortex neurons.* Nature, 1997. **389**(6653): p. 845-848.
40. Graziano, M.S., G.S. Yap, and C.G. Gross, *Coding of visual space by premotor neurons.* Science, 1994. **266**(5187): p. 1054-1057.
41. Pesaran, B., M.J. Nelson, and R.A. Andersen, *Dorsal premotor neurons encode the relative position of the hand, eye, and goal during reach planning.* Neuron, 2006. **51**(1): p. 125-134.
42. Colby, C.L., J.R. Duhamel, and M.E. Goldberg, *Ventral intraparietal area of the macaque: anatomic location and visual response properties.* Journal of Neurophysiology, 1993. **69**(3): p. 902-914.
43. Sakata, H., et al., *Neural Mechanisms of Visual Guidance of Hand Action in the Parietal Cortex of the Monkey.* Cerebral Cortex, 1995.
44. Murata, A., et al., *Selectivity for the shape, size, and orientation of objects for grasping in neurons of monkey parietal area AIP.* Journal of Neurophysiology, 2000. **83**(5): p. 2580-2601.

45. Dum, R.P. and P.L. Strick, *Frontal lobe inputs to the digit representations of the motor areas on the lateral surface of the hemisphere.* The Journal of Neuroscience, 2005. **25**(6): p. 1375-1386.
46. Tomassini, V., et al., *Diffusion-weighted imaging tractography-based parcellation of the human lateral premotor cortex identifies dorsal and ventral subregions with anatomical and functional specializations.* The Journal of Neuroscience, 2007. **27**(38): p. 10259-10269.
47. Hoshi, E. and J. Tanji, *Contrasting neuronal activity in the dorsal and ventral premotor areas during preparation to reach.* Journal of Neurophysiology, 2002. **87**(2): p. 1123-1128.
48. Hoshi, E. and J. Tanji, *Differential involvement of neurons in the dorsal and ventral premotor cortex during processing of visual signals for action planning.* Journal of Neurophysiology, 2006. **95**(6): p. 3596-3616.
49. Hoshi, E. and J. Tanji, *Distinctions between dorsal and ventral premotor areas: anatomical connectivity and functional properties.* Current Opinion in Neurobiology, 2007. **17**(2): p. 234-242.
50. Nakayama, Y., et al., *Transformation of a virtual action plan into a motor plan in the premotor cortex.* The Journal of Neuroscience, 2008. **28**(41): p. 10287-10297.
51. Glaser, J.I., et al., *Population coding of conditional probability distributions in dorsal premotor cortex.* Nature Communications, 2018. **9**(1): p. 1788.
52. O'Shea, J., et al., *Functional specificity of human premotor-motor cortical interactions during action selection.* The European Journal of Neuroscience, 2007. **26**(7): p. 2085-2095.
53. Halsband, U. and R.E. Passingham, *Premotor cortex and the conditions for movement in monkeys (Macaca fascicularis).* Behavioural Brain Research, 1985. **18**(3): p. 269-277.
54. Mitz, A.R., M. Godschalk, and S.P. Wise, *Learning-dependent neuronal activity in the premotor cortex: activity during the acquisition of conditional motor associations.* The Journal of Neuroscience, 1991. **11**(6): p. 1855-1872.
55. Hardwick, R.M., et al., *A quantitative meta-analysis and review of motor learning in the human brain.* Neuroimage, 2013. **67**: p. 283-297.
56. Kornysheva, K. and J. Diedrichsen, *Human premotor areas parse sequences into their spatial and temporal features.* eLife, 2014. **3**: p. e03043.
57. Wiestler, T. and J. Diedrichsen, *Skill learning strengthens cortical representations of motor sequences.* eLife, 2013. **2**: p. e00801.
58. Nachev, P., C. Kennard, and M. Husain, *Functional role of the supplementary and pre-supplementary motor areas.* Nature Reviews. Neuroscience, 2008. **9**(11): p. 856-869.
59. Xu, B., et al., *PreSMA stimulation changes task-free functional connectivity in the fronto-basal-ganglia that correlates with response inhibition efficiency.* Hum Brain Mapp, 2016. **37**(9): p. 3236-49.
60. Penfield, W., *The supplementary motor area of the cerebral cortex.* A.M.A. Archives of Neurology & Psychiatry, 1951. **66**(3): p. 289.
61. Della Sala, S., C. Marchetti, and H. Spinnler, *Right-sided anarchic (alien) hand: a longitudinal study.* Neuropsychologia, 1991. **29**(11): p. 1113-27.
62. Gaymard, B., C. Pierrot-Deseilligny, and S. Rivaud, *Impairment of sequences of memory-guided saccades after supplementary motor area lesions.* Annals of Neurology, 1990. **28**(5): p. 622-626.
63. Gerloff, C., et al., *Stimulation over the human supplementary motor area interferes with the organization of future elements in complex motor sequences.* Brain: A Journal of Neurology, 1997. **120 (Pt 9)**: p. 1587-1602.
64. Luppino, G., et al., *Corticocortical connections of area F3 (SMA-proper) and area F6 (pre-SMA) in the macaque monkey.* The Journal of Comparative Neurology, 1993. **338**(1): p. 114-140.

65. Shima, K. and J. Tanji, *Neuronal activity in the supplementary and presupplementary motor areas for temporal organization of multiple movements.* Journal of Neurophysiology, 2000. **84**(4): p. 2148-2160.
66. Nakamura, K., K. Sakai, and O. Hikosaka, *Neuronal activity in medial frontal cortex during learning of sequential procedures.* Journal of Neurophysiology, 1998. **80**(5): p. 2671-2687.
67. Hikosaka, O., et al., *Activation of human presupplementary motor area in learning of sequential procedures: a functional MRI study.* Journal of Neurophysiology, 1996. **76**(1): p. 617-621.
68. Haith, A.M. and J.W. Krakauer, *Model-based and model-free mechanisms of human motor learning.* Advances in Experimental Medicine and Biology, 2013. **782**: p. 1-21.
69. Nelson, A.B. and A.C. Kreitzer, *Reassessing models of basal ganglia function and dysfunction.* Annual Review of Neuroscience, 2014. **37**(1): p. 117-135.
70. Redgrave, P., et al., *Goal-directed and habitual control in the basal ganglia: implications for Parkinson's disease.* Nature Reviews. Neuroscience, 2010. **11**(11): p. 760-772.
71. Albin, R.L., A.B. Young, and J.B. Penney, *The functional anatomy of disorders of the basal ganglia.* Trends in Neurosciences, 1995. **18**(2): p. 63-64.
72. Parent, A. and L.N. Hazrati, *Functional anatomy of the basal ganglia. I. The cortico-basal ganglia-thalamo-cortical loop.* Brain Research. Brain Research Reviews, 1995. **20**(1): p. 91-127.
73. Nambu, A., H. Tokuno, and M. Takada, *Functional significance of the cortico-subthalamo-pallidal 'hyperdirect' pathway.* Neuroscience Research, 2002. **43**(2): p. 111-117.
74. Nambu, A., *Seven problems on the basal ganglia.* Current Opinion in Neurobiology, 2008. **18**(6): p. 595-604.
75. Maia, T.V. and M.J. Frank, *From reinforcement learning models to psychiatric and neurological disorders.* Nature Neuroscience, 2011. **14**(2): p. 154-162.
76. Frank, M.J., *Dynamic Dopamine Modulation in the Basal Ganglia: A Neurocomputational Account of Cognitive Deficits in Medicated and Nonmedicated Parkinsonism.* Journal of Cognitive Neuroscience, 2005. **17**(1): p. 51-72.
77. Maia, T.V., *Reinforcement learning, conditioning, and the brain: Successes and challenges.* Cognitive, Affective, & Behavioral Neuroscience, 2009. **9**(4): p. 343-364.
78. Maia, T.V., *Two-factor theory, the actor-critic model, and conditioned avoidance.* Learning & Behavior, 2010. **38**(1): p. 50-67.
79. O'Doherty, J., et al., *Dissociable roles of ventral and dorsal striatum in instrumental conditioning.* Science, 2004. **304**(5669): p. 452-4.
80. Mink, J.W., *The basal ganglia: focused selection and inhibition of competing motor programs.* Prog Neurobiol, 1996. **50**(4): p. 381-425.
81. Schultz, W., P. Dayan, and P.R. Montague, *A neural substrate of prediction and reward.* Science, 1997. **275**(5306): p. 1593-1599.
82. Schultz, W., *Reward prediction error.* Current Biology, 2017. **27**(10): p. R369-R371.
83. Jog, M.S., et al., *Building neural representations of habits.* Science, 1999. **286**(5445): p. 1745-1749.
84. Balleine, B.W., M.R. Delgado, and O. Hikosaka, *The role of the dorsal striatum in reward and decision-making.* The Journal of Neuroscience, 2007. **27**(31): p. 8161-8165.
85. Wolpert, D.M., R.C. Miall, and M. Kawato, *Internal models in the cerebellum.* Trends in Cognitive Sciences, 1998. **2**(9): p. 338-347.
86. Ito, M., *Control of mental activities by internal models in the cerebellum.* Nature Reviews. Neuroscience, 2008. **9**(4): p. 304-313.

87. Caligiore, D., et al., *Consensus Paper: Towards a Systems-Level View of Cerebellar Function: the Interplay Between Cerebellum, Basal Ganglia, and Cortex.* Cerebellum, 2017. **16**(1): p. 203-229.
88. Kelly, R.M. and P.L. Strick, *Cerebellar loops with motor cortex and prefrontal cortex of a nonhuman primate.* The Journal of Neuroscience, 2003. **23**(23): p. 8432-8444.
89. Buckner, R.L., et al., *The organization of the human cerebellum estimated by intrinsic functional connectivity.* Journal of Neurophysiology, 2011. **106**(5): p. 2322-2345.
90. Bostan, A.C. and P.L. Strick, *The basal ganglia and the cerebellum: nodes in an integrated network.* Nat Rev Neurosci, 2018. **19**(6): p. 338-350.
91. Sterling, P. and S. Laughlin, *Principles of neural design*. 2015: The MIT Press.
92. Morton, S.M. and A.J. Bastian, *Cerebellar contributions to locomotor adaptations during splitbelt treadmill walking.* The Journal of Neuroscience, 2006. **26**(36): p. 9107-9116.
93. Martin, T.A., et al., *Throwing while looking through prisms. I. Focal olivocerebellar lesions impair adaptation.* Brain, 1996. **119 (Pt 4)**: p. 1183-98.
94. Smith, M.A. and R. Shadmehr, *Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration.* Journal of Neurophysiology, 2005. **93**(5): p. 2809-2821.
95. Hebb, D.O., *The organization of behavior: A neuropsychological theory*. 1949, Mahwah, N.J: L. Erlbaum Associates.
96. Davis, G.W. and C.S. Goodman, *Synapse-specific control of synaptic efficacy at the terminals of a single neuron.* Nature, 1998. **392**(6671): p. 82-86.
97. O'Brien, R.J., et al., *Activity-dependent modulation of synaptic AMPA receptor accumulation.* Neuron, 1998. **21**(5): p. 1067-1078.
98. Turrigiano, G.G., et al., *Activity-dependent scaling of quantal amplitude in neocortical neurons.* Nature, 1998. **391**(6670): p. 892-896.
99. Caporale, N. and Y. Dan, *Spike timing-dependent plasticity: a Hebbian learning rule.* Annual Review of Neuroscience, 2008. **31**: p. 25-46.
100. Bliss, T.V. and T. Lomo, *Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path.* The Journal of Physiology, 1973. **232**(2): p. 331-356.
101. Stanton, P.K. and T.J. Sejnowski, *Associative long-term depression in the hippocampus induced by hebbian covariance.* Nature, 1989. **339**(6221): p. 215-218.
102. Bliss, T.V. and G.L. Collingridge, *A synaptic model of memory: long-term potentiation in the hippocampus.* Nature, 1993. **361**(6407): p. 31-39.
103. Engert, F. and T. Bonhoeffer, *Dendritic spine changes associated with hippocampal long-term synaptic plasticity.* Nature, 1999. **399**(6731): p. 66-70.
104. Matsuzaki, M., et al., *Structural basis of long-term potentiation in single dendritic spines.* Nature, 2004. **429**(6993): p. 761-766.
105. Harvey, C.D. and K. Svoboda, *Locally dynamic synaptic learning rules in pyramidal neuron dendrites.* Nature, 2007. **450**(7173): p. 1195-1200.
106. Artola, A. and W. Singer, *Long-term potentiation and NMDA receptors in rat visual cortex.* Nature, 1987. **330**(6149): p. 649-652.
107. Iriki, A., et al., *Long-term potentiation in the motor cortex.* Science, 1989. **245**(4924): p. 1385-1387.
108. Liu, Q.-s., L. Pu, and M.-m. Poo, *Repeated cocaine exposure in vivo facilitates LTP induction in midbrain dopamine neurons.* Nature, 2005. **437**(7061): p. 1027-1031.
109. Xu, T., et al., *Rapid formation and selective stabilization of synapses for enduring motor memories.* Nature, 2009. **462**(7275): p. 915-919.
110. Yang, G., F. Pan, and W.-B. Gan, *Stably maintained dendritic spines are associated with lifelong memories.* Nature, 2009. **462**(7275): p. 920-924.
111. Bejjanki, V.R., et al., *Perceptual learning as improved probabilistic inference in early sensory areas.* Nature Neuroscience, 2011. **14**(5): p. 642-648.

112. Seriès, P., P.E. Latham, and A. Pouget, *Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations.* Nature Neuroscience, 2004. **7**(10): p. 1129-1135.
113. Makino, H., et al., *Circuit mechanisms of sensorimotor learning.* Neuron, 2016. **92**(4): p. 705-721.
114. Athalye, V.R., et al., *Emergence of coordinated neural dynamics underlies neuroprosthetic learning and skillful control.* Neuron, 2017. **93**(4): p. 955-970.e5.
115. Kleim, J.A., S. Barbay, and R.J. Nudo, *Functional reorganization of the rat motor cortex following motor skill learning.* Journal of Neurophysiology, 1998. **80**(6): p. 3321-3325.
116. Pascual-Leone, A., J. Grafman, and M. Hallett, *Modulation of cortical motor output maps during development of implicit and explicit knowledge.* Science, 1994. **263**(5151): p. 1287-1289.
117. Pascual-Leone, A., et al., *Modulation of muscle responses evoked by transcranial magnetic stimulation during the acquisition of new fine motor skills.* Journal of Neurophysiology, 1995. **74**(3): p. 1037-1045.
118. Hlustík, P., et al., *Cortical plasticity during three-week motor skill learning.* Journal of Clinical Neurophysiology, 2004. **21**(3): p. 180-191.
119. Jacobs, K.M. and J.P. Donoghue, *Reshaping the cortical motor map by unmasking latent intracortical connections.* Science, 1991. **251**(4996): p. 944-947.
120. Kolasinski, J., et al., *The dynamics of cortical GABA in human motor learning.* bioRxiv, 2018.
121. Alexander, G.E. and M.D. Crutcher, *Functional architecture of basal ganglia circuits: neural substrates of parallel processing.* Trends in Neurosciences, 1990. **13**(7): p. 266-271.
122. da Silva, J.A., et al., *Dopamine neuron activity before action initiation gates and invigorates future movements.* Nature, 2018. **554**(7691): p. 244-248.
123. Graybiel, A.M., *The basal ganglia: learning new tricks and loving it.* Current Opinion in Neurobiology, 2005. **15**(6): p. 638-644.
124. Hosp, J.A., et al., *Dopaminergic projections from midbrain to primary motor cortex mediate motor skill learning.* The Journal of Neuroscience, 2011. **31**(7): p. 2481-2487.
125. Guo, L., et al., *Dynamic rewiring of neural circuits in the motor cortex in mouse models of Parkinson's disease.* Nature Neuroscience, 2015. **18**(9): p. 1299-1309.
126. Ramkumar, P., et al., *Premotor and motor cortices encode reward.* Plos One, 2016. **11**(8): p. e0160851.
127. Santos, F.J., et al., *Corticostriatal dynamics encode the refinement of specific behavioral variability during skill learning.* eLife, 2015. **4**: p. e09423.
128. Yang, Y. and S.G. Lisberger, *Purkinje-cell plasticity and cerebellar motor learning are graded by complex-spike duration.* Nature, 2014. **510**(7506): p. 529-32.
129. Chen, S.X., et al., *Subtype-specific plasticity of inhibitory circuits in motor cortex during motor learning.* Nature Neuroscience, 2015. **18**(8): p. 1109-1115.
130. Picard, N., Y. Matsuzaka, and P.L. Strick, *Extended practice of a motor skill is associated with reduced metabolic activity in M1.* Nature Neuroscience, 2013. **16**(9): p. 1340-1347.
131. Penhune, V.B. and J. Doyon, *Dynamic cortical and subcortical networks in learning and delayed recall of timed motor sequences.* The Journal of Neuroscience, 2002. **22**(4): p. 1397-1406.
132. Steele, C.J. and V.B. Penhune, *Specific increases within global decreases: a functional magnetic resonance imaging investigation of five days of motor sequence learning.* The Journal of Neuroscience, 2010. **30**(24): p. 8332-8341.
133. Ma, L., et al., *Changes in regional activity are accompanied with changes in inter-regional connectivity during 4 weeks motor learning.* Brain Research, 2010. **1318**: p. 64-76.

134. Wymbs, N.F. and S.T. Grafton, *The Human Motor System Supports Sequence-Specific Representations over Multiple Training-Dependent Timescales.* Cerebral Cortex, 2015. **25**(11): p. 4213-4225.
135. Peters, A.J., et al., *Reorganization of corticospinal output during motor learning.* Nature Neuroscience, 2017. **20**(8): p. 1133-1141.
136. Han, Y., et al., *Gray matter density and white matter integrity in pianists' brain: a combined structural and diffusion tensor MRI study.* Neuroscience Letters, 2009. **459**(1): p. 3-6.
137. Gaser, C. and G. Schlaug, *Brain structures differ between musicians and non-musicians.* The Journal of Neuroscience, 2003. **23**(27): p. 9240-9245.
138. Sadtler, P.T., et al., *Neural constraints on learning.* Nature, 2014. **512**(7515): p. 423-6.
139. Golub, M.D., et al., *Learning by neural reassociation.* Nature Neuroscience, 2018. **21**(4): p. 607-616.
140. Perich, M.G., J.A. Gallego, and L. Miller, *A neural population mechanism for rapid learning.* BioRxiv, 2017.
141. Makino, H., et al., *Transformation of Cortex-wide Emergent Properties during Motor Learning.* Neuron, 2017. **94**(4): p. 880-890.e8.
142. Bassett, D.S., et al., *Dynamic reconfiguration of human brain networks during learning.* Proc Natl Acad Sci U S A, 2011. **108**(18): p. 7641-6.
143. Bassett, D.S., et al., *Learning-induced autonomy of sensorimotor systems.* Nature Neuroscience, 2015. **18**(5): p. 744-751.
144. Albert, N.B., E.M. Robertson, and R.C. Miall, *The resting human brain and motor learning.* Curr Biol, 2009. **19**(12): p. 1023-7.
145. Sami, S. and R.C. Miall, *Graph network analysis of immediate motor-learning induced changes in resting state BOLD.* Frontiers in Human Neuroscience, 2013. **7**: p. 166.
146. Taubert, M., et al., *Long-term effects of motor training on resting-state networks and underlying brain structure.* Neuroimage, 2011. **57**(4): p. 1492-1498.
147. Vahdat, S., et al., *Functionally specific changes in resting-state sensorimotor networks after motor learning.* J Neurosci, 2011. **31**(47): p. 16907-15.
148. Sampaio-Baptista, C., et al., *Changes in functional connectivity and GABA levels with long-term motor learning.* Neuroimage, 2015. **106**: p. 15-20.
149. Ma, L., et al., *Changes occur in resting state network of motor system during 4 weeks of motor skill learning.* Neuroimage, 2011. **58**(1): p. 226-233.
150. Sami, S., E.M. Robertson, and R.C. Miall, *The time course of task-specific memory consolidation effects in resting state networks.* J Neurosci, 2014. **34**(11): p. 3982-92.
151. Galea, J.M., et al., *The dissociable effects of punishment and reward on motor learning.* Nat Neurosci, 2015. **18**(4): p. 597-602.
152. Wachter, T., et al., *Differential effect of reward and punishment on procedural learning.* J Neurosci, 2009. **29**(2): p. 436-43.
153. Stagg, C.J., et al., *Local GABA concentration is related to network-level resting functional connectivity.* eLife, 2014. **3**: p. e01465.
154. Kapogiannis, D., et al., *Posteromedial cortex glutamate and GABA predict intrinsic functional connectivity of the default mode network.* Neuroimage, 2013. **64**: p. 112-119.
155. Nissen, M.J. and P. Bullemer, *Attentional requirements of learning: Evidence from performance measures.* Cognitive Psychology, 1987. **19**(1): p. 1-32.
156. Willingham, D.B., M.J. Nissen, and P. Bullemer, *On the development of procedural knowledge.* J Exp Psychol Learn Mem Cogn, 1989. **15**(6): p. 1047-60.
157. Robertson, E.M., *The serial reaction time task: implicit motor skill learning?* J Neurosci, 2007. **27**(38): p. 10073-5.

158. Corkin, S., *Acquisition of motor skill after bilateral medial temporal-lobe excision.* Neuropsychologia, 1968. **6**(3): p. 255-265.
159. Dayan, E. and L.G. Cohen, *Neuroplasticity subserving motor skill learning.* Neuron, 2011. **72**(3): p. 443-54.
160. Robertson, E.M., *Memory instability as a gateway to generalization.* PLoS Biol, 2018. **16**(3): p. e2004633.
161. Chen, X., P. Holland, and J.M. Galea, *The effects of reward and punishment on motor skill learning.* Current Opinion in Behavioral Sciences, 2018. **20**: p. 83-88.
162. Damadian, R., M. Goldsmith, and L. Minkoff, *NMR in cancer: XVI. FONAR image of the live human body.* Physiol Chem Phys, 1977. **9**(1): p. 97-100, 108.
163. Pauling, L. and C.D. Coryell, *The Magnetic Properties and Structure of Hemoglobin, Oxyhemoglobin and Carbonmonoxyhemoglobin.* Proc Natl Acad Sci U S A, 1936. **22**(4): p. 210-6.
164. Jezzard, P., P.M. Matthews, and S.M. Smith, *Functional MRI: an introduction to methods*. 2001, Oxford: Oxford University Press.
165. Ogawa, S., et al., *Brain magnetic resonance imaging with contrast dependent on blood oxygenation.* Proc Natl Acad Sci U S A, 1990. **87**(24): p. 9868-72.
166. Stehling, M.K., et al., *Inversion-recovery echo-planar imaging (IR-EPI) at 0.5 T.* Magn Reson Med, 1990. **13**(3): p. 514-7.
167. Mansfield, P. and A.A. Maudsley, *Medical imaging by NMR.* Br J Radiol, 1977. **50**(591): p. 188-94.
168. Kwong, K.K., et al., *Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation.* Proc Natl Acad Sci U S A, 1992. **89**(12): p. 5675-9.
169. Bandettini, P.A., et al., *Time course EPI of human brain function during task activation.* Magn Reson Med, 1992. **25**(2): p. 390-7.
170. Ogawa, S., et al., *Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging.* Proc Natl Acad Sci U S A, 1992. **89**(13): p. 5951-5.
171. Kanwisher, N., J. McDermott, and M.M. Chun, *The fusiform face area: a module in human extrastriate cortex specialized for face perception.* J Neurosci, 1997. **17**(11): p. 4302-11.
172. Biswal, B., et al., *Functional connectivity in the motor cortex of resting human brain using echo-planar MRI.* Magn Reson Med, 1995. **34**(4): p. 537-41.
173. Power, J.D., et al., *Functional network organization of the human brain.* Neuron, 2011. **72**(4): p. 665-78.
174. Lewis, C.M., et al., *Learning sculpts the spontaneous activity of the resting human brain.* Proc Natl Acad Sci U S A, 2009. **106**(41): p. 17558-63.
175. Gotts, S.J., et al., *Fractionation of social brain circuits in autism spectrum disorders.* Brain, 2012. **135**(Pt 9): p. 2711-25.
176. Agosta, F., et al., *Resting state fMRI in Alzheimer's disease: beyond the default mode network.* Neurobiol Aging, 2012. **33**(8): p. 1564-78.
177. Winder, A.T., et al., *Weak correlations between hemodynamic signals and ongoing neural activity during the resting state.* Nat Neurosci, 2017. **20**(12): p. 1761-1769.
178. Ma, Y., et al., *Resting-state hemodynamics are spatiotemporally coupled to synchronized and symmetric neural activity in excitatory neurons.* Proc Natl Acad Sci U S A, 2016. **113**(52): p. E8463-E8471.
179. Murayama, Y., et al., *Relationship between neural and hemodynamic signals during spontaneous activity studied with temporal kernel CCA.* Magn Reson Imaging, 2010. **28**(8): p. 1095-103.
180. Fries, P., et al., *Modulation of oscillatory neuronal synchronization by selective visual attention.* Science, 2001. **291**(5508): p. 1560-3.
181. Schoffelen, J.M., R. Oostenveld, and P. Fries, *Neuronal coherence as a mechanism of effective corticospinal interaction.* Science, 2005. **308**(5718): p. 111-3.

182. Gray, C.M., et al., *Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties.* Nature, 1989. **338**(6213): p. 334-7.
183. Power, J.D., et al., *Methods to detect, characterize, and remove motion artifact in resting state fMRI.* Neuroimage, 2014. **84**: p. 320-41.
184. Power, J.D., B.L. Schlaggar, and S.E. Petersen, *Recent progress and outstanding issues in motion correction in resting state fMRI.* Neuroimage, 2015. **105**: p. 536-51.
185. Power, J.D., et al., *Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion.* Neuroimage, 2012. **59**(3): p. 2142-54.
186. Siegel, J.S., et al., *Statistical improvements in functional magnetic resonance imaging analyses produced by censoring high-motion data points.* Hum Brain Mapp, 2014. **35**(5): p. 1981-96.
187. White, N., et al., *PROMO: Real-time prospective motion correction in MRI using image-based tracking.* Magnetic Resonance in Medicine, 2010. **63**(1): p. 91-105.
188. Birn, R.M., et al., *The influence of physiological noise correction on test-retest reliability of resting-state functional connectivity.* Brain Connect, 2014. **4**(7): p. 511-22.
189. Murphy, K., R.M. Birn, and P.A. Bandettini, *Resting-state fMRI confounds and cleanup.* Neuroimage, 2013. **80**: p. 349-59.
190. Biswal, B., E.A. DeYoe, and J.S. Hyde, *Reduction of physiological fluctuations in fMRI using digital filters.* Magn Reson Med, 1996. **35**(1): p. 107-13.
191. Glover, G.H., T.Q. Li, and D. Ress, *Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR.* Magn Reson Med, 2000. **44**(1): p. 162-7.
192. Jo, H.J., et al., *Effective Preprocessing Procedures Virtually Eliminate Distance-Dependent Motion Artifacts in Resting State FMRI.* J Appl Math, 2013. **2013**.
193. Jo, H.J., et al., *Mapping sources of correlation in resting state FMRI, with artifact detection and removal.* Neuroimage, 2010. **52**(2): p. 571-82.
194. Salimi-Khorshidi, G., et al., *Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers.* Neuroimage, 2014. **90**: p. 449-68.
195. Kundu, P., et al., *Differentiating BOLD and non-BOLD signals in fMRI time series using multi-echo EPI.* Neuroimage, 2012. **60**(3): p. 1759-70.
196. Posse, S., et al., *Enhancement of BOLD-contrast sensitivity by single-shot multi-echo functional MR imaging.* Magn Reson Med, 1999. **42**(1): p. 87-97.
197. Evans, J.W., et al., *Separating slow BOLD from non-BOLD baseline drifts using multi-echo fMRI.* Neuroimage, 2015. **105**: p. 189-97.
198. Cox, R.W., *AFNI: software for analysis and visualization of functional magnetic resonance neuroimages.* Comput Biomed Res, 1996. **29**(3): p. 162-73.
199. Jenkinson, M., et al., *Fsl.* Neuroimage, 2012. **62**(2): p. 782-90.
200. Sladky, R., et al., *Slice-timing effects and their correction in functional MRI.* NeuroImage, 2011. **58**(2): p. 588-594.
201. Griffanti, L., et al., *ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging.* Neuroimage, 2014. **95**: p. 232-47.
202. Bassett, D.S. and M.G. Mattar, *A Network Neuroscience of Human Learning: Potential to Inform Quantitative Theories of Brain and Behavior.* Trends Cogn Sci, 2017. **21**(4): p. 250-264.
203. Smith, S.M., et al., *Temporally-independent functional modes of spontaneous brain activity.* Proc Natl Acad Sci U S A, 2012. **109**(8): p. 3131-6.
204. Steel, A., et al., *Metabolite-cycled density-weighted concentric rings k-space trajectory (DW-CRT) enables high-resolution 1 H magnetic resonance spectroscopic imaging at 3-Tesla.* Sci Rep, 2018. **8**(1): p. 7792.

205. Shulman, R.G., et al., *Cellular applications of 31P and 13C nuclear magnetic resonance.* Science, 1979. **205**(4402): p. 160-6.
206. Hahn, E.L. and D.E. Maxwell, *Spin Echo Measurements of Nuclear Spin Coupling in Molecules.* Physical Review, 1952. **88**(5): p. 1070-1084.
207. Bottomley, P.A., et al., *Spatial localization in 31P and 13C NMR spectroscopy in Vivo using surface coils.* Magnetic Resonance in Medicine, 1984. **1**(3): p. 410-413.
208. Frahm, J., K.-D. Merboldt, and W. Hänicke, *Localized proton spectroscopy using stimulated echoes.* Journal of Magnetic Resonance (1969), 1987. **72**(3): p. 502-508.
209. Scheenen, T.W.J., et al., *Short echo time 1H-MRSI of the human brain at 3T with minimal chemical shift displacement errors using adiabatic refocusing pulses.* Magnetic Resonance in Medicine, 2007. **59**(1): p. 1-6.
210. Chiew, M., et al., *Density-weighted concentric rings k-space trajectory for 1H magnetic resonance spectroscopic imaging at 7 T.* NMR in Biomedicine, 2017. **31**(1): p. e3838.
211. Emir, U.E., et al., *Non-water-suppressed short-echo-time magnetic resonance spectroscopic imaging using a concentric ring k-space trajectory.* NMR in Biomedicine, 2017. **30**(7): p. e3714.
212. Donadieu, M., et al., *Metabolic voxel-based analysis of the complete human brain using fast 3D-MRSI: Proof of concept in multiple sclerosis.* Journal of Magnetic Resonance Imaging, 2016. **44**(2): p. 411-419.
213. Otazo, R., et al., *Signal-to-noise ratio and spectral linewidth improvements between 1.5 and 7 Tesla in proton echo-planar spectroscopic imaging.* Magnetic Resonance in Medicine, 2006. **56**(6): p. 1200-1210.
214. Furuyama, J.K., N.E. Wilson, and M.A. Thomas, *Spectroscopic imaging using concentrically circular echo-planar trajectories in vivo.* Magnetic Resonance in Medicine, 2011. **67**(6): p. 1515-1522.
215. Zhu, H. and P.B. Barker, *MR Spectroscopy and Spectroscopic Imaging of the Brain.* Methods in molecular biology (Clifton, N.J.), 2011. **711**: p. 203-226.
216. Hock, A., et al., *Non-water-suppressed proton MR spectroscopy improves spectral quality in the human spinal cord.* Magnetic Resonance in Medicine, 2012. **69**(5): p. 1253-1260.
217. Pijnappel, W.W.F., et al., *SVD-based quantification of magnetic resonance signals.* Journal of Magnetic Resonance (1969), 1992. **97**(1): p. 122-134.
218. Cabanes, E., et al., *Optimization of Residual Water Signal Removal by HLSVD on Simulated Short Echo Time Proton MR Spectra of the Human Brain.* Journal of Magnetic Resonance, 2001. **150**(2): p. 116-125.
219. Bilgic, B., et al., *Fast quantitative susceptibility mapping with L1-regularization and automatic parameter selection.* Magnetic Resonance in Medicine, 2013. **72**(5): p. 1444-1459.
220. Provencher, S.W., *Automatic quantitation of localized in vivo 1H spectra with LCModel.* NMR Biomed, 2001. **14**(4): p. 260-4.
221. Cavassila, S., et al., *Cramér–Rao bounds: an evaluation tool for quantitation.* NMR in Biomedicine, 2001. **14**(4): p. 278-283.
222. Kreis, R., *The trouble with quality filtering based on relative Cramer-Rao lower bounds.* Magn Reson Med, 2016. **75**(1): p. 15-8.
223. Mescher, M., et al., *Simultaneous in vivo spectral editing and water suppression.* NMR in Biomedicine, 1998. **11**(6): p. 266-272.
224. Stagg, C.J., et al., *Relationship between physiological measures of excitability and levels of glutamate and GABA in the human motor cortex.* J Physiol, 2011. **589**(Pt 23): p. 5845-55.
225. Kolasinski, J., et al., *A Mechanistic Link from GABA to Cortical Architecture and Perception.* Curr Biol, 2017. **27**(11): p. 1685-1691 e3.

226. Albouy, G., et al., *Hippocampus and striatum: dynamics and interaction during acquisition and sleep-related motor sequence memory consolidation.* Hippocampus, 2013. **23**(11): p. 985-1004.
227. Albouy, G., et al., *Interaction between hippocampal and striatal systems predicts subsequent consolidation of motor sequence memory.* Plos One, 2013. **8**(3): p. e59490.
228. Thomas, C. and C.I. Baker, *Teaching an adult brain new tricks: a critical review of evidence for training-dependent structural plasticity in humans.* Neuroimage, 2013. **73**: p. 225-236.
229. Trefler, A., et al., *Impact of time-of-day on brain morphometric measures derived from T1-weighted magnetic resonance imaging.* Neuroimage, 2016. **133**: p. 41-52.
230. Heitger, M.H., et al., *Motor learning-induced changes in functional brain connectivity as revealed by means of graph-theoretical network analysis.* Neuroimage, 2012. **61**(3): p. 633-650.
231. Schlaffke, L., et al., *Dynamic changes of resting state connectivity related to the acquisition of a lexico-semantic skill.* Neuroimage, 2017. **146**: p. 429-437.
232. Shannon, B.J., et al., *Morning-evening variation in human brain metabolism and memory circuits.* Journal of Neurophysiology, 2013. **109**(5): p. 1444-1456.
233. Sagi, Y., et al., *Learning in the fast lane: new insights into neuroplasticity.* Neuron, 2012. **73**(6): p. 1195-1203.
234. Peirce, J.W., *PsychoPy--Psychophysics software in Python.* Journal of Neuroscience Methods, 2007. **162**(1-2): p. 8-13.
235. Bates, D., et al., *Fitting Linear Mixed-Effects Models Using lme4.* 2015, 2015. **67**(1): p. 48.
236. Fischl, B., et al., *Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain.* Neuron, 2002. **33**(3): p. 341-55.
237. Gonzalez-Castillo, J., et al., *The spatial structure of resting state connectivity stability on the scale of minutes.* Frontiers in Neuroscience, 2014. **8**: p. 138.
238. Saad, Z.S., et al., *Correcting brain-wide correlation differences in resting-state FMRI.* Brain connectivity, 2013. **3**(4): p. 339-352.
239. Saad, Z.S., et al., *Trouble at rest: how correlation patterns and group differences become distorted after global signal regression.* Brain connectivity, 2012. **2**(1): p. 25-32.
240. Chen, G., et al., *Linear mixed-effects modeling approach to FMRI group analysis.* Neuroimage, 2013. **73**: p. 176-90.
241. Kriegeskorte, N., et al., *Circular analysis in systems neuroscience: the dangers of double dipping.* Nature Neuroscience, 2009. **12**(5): p. 535-540.
242. Pinherio, J., et al., *nlme: Linear and Nonlinear Mixed Effects Models.* 2018.
243. Lenth, R.V., *Least-Squares Means: The R Package lsmeans.* 2016, 2016. **69**(1): p. 33.
244. Baldassano, C., et al., *Two Distinct Scene-Processing Networks Connecting Vision and Memory.* eNeuro, 2016. **3**(5).
245. Silson, E.H., et al., *A Retinotopic Basis for the Division of High-Level Scene Processing between Lateral and Ventral Human Occipitotemporal Cortex.* The Journal of Neuroscience, 2015. **35**(34): p. 11921-11935.
246. Silson, E.H., A.D. Steel, and C.I. Baker, *Scene-Selectivity and Retinotopy in Medial Parietal Cortex.* Frontiers in Human Neuroscience, 2016. **10**: p. 412.
247. Keller, T.A. and M.A. Just, *Structural and functional neuroplasticity in human learning of spatial routes.* Neuroimage, 2016. **125**: p. 256-266.
248. Doyon, J., et al., *Current issues related to motor sequence learning in humans.* Current opinion in behavioral sciences, 2018. **20**: p. 89-97.
249. Doyon, J., V. Penhune, and L.G. Ungerleider, *Distinct contribution of the cortico-striatal and cortico-cerebellar systems to motor skill learning.* Neuropsychologia, 2003. **41**(3): p. 252-262.

250. Miyachi, S., O. Hikosaka, and X. Lu, *Differential activation of monkey striatal neurons in the early and late stages of procedural learning.* Experimental Brain Research, 2002. **146**(1): p. 122-126.
251. Seger, C.A., *How do the basal ganglia contribute to categorization? Their roles in generalization, response selection, and learning via feedback.* Neuroscience and Biobehavioral Reviews, 2008. **32**(2): p. 265-278.
252. Schendan, H.E., et al., *An FMRI study of the role of the medial temporal lobe in implicit and explicit sequence learning.* Neuron, 2003. **37**(6): p. 1013-1025.
253. Schendan, H.E., et al., *Sequence? What Sequence?: the human medial temporal lobe and sequence learning.* Molecular Psychiatry, 2003. **8**(11): p. 896-897.
254. Park, B., et al., *Are brain networks stable during a 24-hour period?* Neuroimage, 2012. **59**(1): p. 456-66.
255. Hodkinson, D.J., et al., *Circadian and homeostatic modulation of functional connectivity and regional cerebral blood flow in humans under normal entrained conditions.* Journal of Cerebral Blood Flow and Metabolism, 2014. **34**(9): p. 1493-1499.
256. Jiang, C., et al., *Diurnal Variations in Neural Activity of Healthy Human Brain Decoded with Resting-State Blood Oxygen Level Dependent fMRI.* Front Hum Neurosci, 2016. **10**: p. 634.
257. Thomas, C., et al., *Impact of time-of-day on diffusivity measures of brain tissue derived from diffusion tensor imaging.* Neuroimage, 2018. **173**: p. 25-34.
258. Chodobski, A. and J. Szmydynger-Chodobska, *Choroid plexus: target for polypeptides and site of their synthesis.* Microsc Res Tech, 2001. **52**(1): p. 65-82.
259. Gregory, M.D., et al., *Thinking About a Task Is Associated with Increased Connectivity in Regions Activated by Task Performance.* Brain Connect, 2016. **6**(2): p. 164-8.
260. Freedberg, M., et al., *Comparing the effects of positive and negative feedback in information-integration category learning.* Mem Cognit, 2017. **45**(1): p. 12-25.
261. Ashby, F.G., J.M. Ennis, and B.J. Spiering, *A neurobiological theory of automaticity in perceptual categorization.* Psychological Review, 2007. **114**(3): p. 632-656.
262. Fouragnan, E., C. Retzler, and M.G. Philiastides, *Separate neural representations of prediction error valence and surprise: Evidence from an fMRI meta-analysis.* Human Brain Mapping, 2018.
263. van der Kooij, K., et al., *Reward abundance interferes with error-based learning in a visuomotor adaptation task.* Plos One, 2018. **13**(3): p. e0193002.
264. van der Kooij, K. and K.E. Overvliet, *Rewarding imperfect motor performance reduces adaptive changes.* Experimental Brain Research, 2016. **234**(6): p. 1441-1450.
265. Hasson, C.J., J. Manczurowsky, and S.-C. Yen, *A reinforcement learning approach to gait training improves retention.* Frontiers in Human Neuroscience, 2015. **9**: p. 459.
266. Quattrocchi, G., et al., *Reward and punishment enhance motor adaptation in stroke.* Journal of Neurology, Neurosurgery, and Psychiatry, 2017.
267. Therrien, A.S., D.M. Wolpert, and A.J. Bastian, *Effective reinforcement learning following cerebellar damage requires a balance between exploration and motor noise.* Brain: A Journal of Neurology, 2016. **139**(Pt 1): p. 101-114.
268. Shmuelof, L., et al., *Overcoming motor "forgetting" through reinforcement of learned actions.* The Journal of Neuroscience, 2012. **32**(42): p. 14617-14621.
269. Pekny, S.E., S.E. Criscimagna-Hemminger, and R. Shadmehr, *Protection and expression of human motor memories.* The Journal of Neuroscience, 2011. **31**(39): p. 13829-13839.
270. Song, S. and L.G. Cohen, *Conscious recall of different aspects of skill memory.* Frontiers in Behavioral Neuroscience, 2014. **8**: p. 233.
271. Tversky, A. and D. Kahneman, *The framing of decisions and the psychology of choice.* Science, 1981. **211**(4481): p. 453-8.

272. Manohar, S.G., et al., *Reward Pays the Cost of Noise Reduction in Motor and Cognitive Control.* Curr Biol, 2015. **25**(13): p. 1707-16.
273. Galea, J.M., et al., *Punishment-induced behavioral and neurophysiological variability reveals dopamine-dependent selection of kinematic movement parameters.* J Neurosci, 2013. **33**(9): p. 3981-8.
274. Wilkinson, L., et al., *The contribution of primary motor cortex is essential for probabilistic implicit sequence learning: evidence from theta burst magnetic stimulation.* Journal of Cognitive Neuroscience, 2010. **22**(3): p. 427-436.
275. Spencer, R.M., et al., *Disrupted timing of discontinuous but not continuous movements by cerebellar lesions.* Science, 2003. **300**(5624): p. 1437-9.
276. Wilkinson, L., et al., *Online feedback enhances early consolidation of motor sequence learning and reverses recall deficit from transcranial stimulation of motor cortex.* Cortex, 2015. **71**: p. 134-47.
277. Robertson, E.M., A. Pascual-Leone, and D.Z. Press, *Awareness modifies the skill-learning benefits of sleep.* Current Biology, 2004. **14**(3): p. 208-212.
278. Tunovic, S., D.Z. Press, and E.M. Robertson, *A physiological signal that prevents motor skill improvements during consolidation.* The Journal of Neuroscience, 2014. **34**(15): p. 5302-5310.
279. Song, S., et al., *Practice structure improves unconscious transitional memories by increasing synchrony in a premotor network.* J Cogn Neurosci, 2015. **27**(8): p. 1503-12.
280. Witt, J.K., J. Ashe, and D.T. Willingham, *An egocentric frame of reference in implicit motor sequence learning.* Psychol Res, 2008. **72**(5): p. 542-52.
281. Taylor, J.A., J.W. Krakauer, and R.B. Ivry, *Explicit and implicit contributions to learning in a sensorimotor adaptation task.* J Neurosci, 2014. **34**(8): p. 3023-32.
282. Maddox, W.T. and C.J. Bohil, *A theoretical framework for understanding the effects of simultaneous base-rate and payoff manipulations on decision criterion learning in perceptual categorization.* J Exp Psychol Learn Mem Cogn, 2003. **29**(2): p. 307-20.
283. Destrebecqz, A. and A. Cleeremans, *Temporal effects in sequence learning*, in *Attention and implicit learning*, L. Jiménez, Editor. 2003, John Benjamins Publishing Company: Amsterdam. p. 181-213.
284. Destrebecqz, A., et al., *The neural correlates of implicit and explicit sequence learning: Interacting networks revealed by the process dissociation procedure.* Learning & Memory, 2005. **12**(5): p. 480-490.
285. Cohen, A., R.I. Ivry, and S.W. Keele, *Attention and structure in sequence learning.* Journal of Experimental Psychology: Learning, Memory, and Cognition, 1990. **16**(1): p. 17-30.
286. Heuer, H. and V. Schmidtke, *Secondary-task effects on sequence learning.* Psychological research, 1996. **59**(2): p. 119-133.
287. Frensch, P.A., D. Wenke, and D. Rünger, *A secondary tone-counting task suppresses expression of knowledge in the serial reaction task.* Journal of Experimental Psychology: Learning, Memory, and Cognition, 1999. **25**(1): p. 260-274.
288. Rah, S.K., A.S. Reber, and A.T. Hsiao, *Another wrinkle on the dual-task SRT experiment: it's probably not dual task.* Psychonomic Bulletin & Review, 2000. **7**(2): p. 309-313.
289. Jiménez, L. and G.A. Vázquez, *Sequence learning under dual-task conditions: alternatives to a resource-based account.* Psychological research, 2005. **69**(5-6): p. 352-368.
290. Codol, O., P.J. Holland, and J.M. Galea, *The relationship between reinforcement and explicit control during visuomotor adaptation.* Scientific reports, 2018. **8**(1): p. 9121.

291. Holland, P., O. Codol, and J.M. Galea, *Contribution of explicit processes to reinforcement-based motor learning.* Journal of Neurophysiology, 2018. **119**(6): p. 2241-2255.
292. DeCaro, M.S., et al., *Choking under pressure: multiple routes to skill failure.* Journal of Experimental Psychology. General, 2011. **140**(3): p. 390-406.
293. Beilock, S.L. and T.H. Carr, *On the fragility of skilled performance: What governs choking under pressure?* Journal of Experimental Psychology: General, 2001. **130**(4): p. 701-725.
294. Sanders, S. and B. Walia, *Shirking and "choking" under incentive-based pressure: A behavioral economic theory of performance production.* Economics letters, 2012. **116**(3): p. 363-366.
295. Baumeister, R.F. and C.J. Showers, *A review of paradoxical performance effects: Choking under pressure in sports and mental tests.* European Journal of Social Psychology, 1986. **16**(4): p. 361-383.
296. Hall, K.G. and R.A. Magill, *Variability of practice and contextual interference in motor skill learning.* Journal of Motor Behavior, 1995. **27**(4): p. 299-309.
297. Cross, E.S., P.J. Schmitt, and S.T. Grafton, *Neural substrates of contextual interference during motor learning support a model of active preparation.* Journal of Cognitive Neuroscience, 2007. **19**(11): p. 1854-1871.
298. Wulf, G. and R.A. Schmidt, *Variability in Practice.* Journal of Motor Behavior, 1988. **20**(2): p. 133-149.
299. Estes, W.K., *Toward a statistical theory of learning.* Psychological Review, 1950. **57**(2): p. 94-107.
300. Worthy, D.A. and W.T. Maddox, *A Comparison Model of Reinforcement-Learning and Win-Stay-Lose-Shift Decision-Making Processes: A Tribute to W.K. Estes.* J Math Psychol, 2014. **59**: p. 41-49.
301. Manley, H., P. Dayan, and J. Diedrichsen, *When money is not enough: awareness, success, and variability in motor learning.* Plos One, 2014. **9**(1): p. e86580.
302. den Ouden, H.E.M., et al., *Dissociable effects of dopamine and serotonin on reversal learning.* Neuron, 2013. **80**(4): p. 1090-1100.
303. Robertson, E.M., et al., *The role of the dorsolateral prefrontal cortex during sequence learning is specific for spatial information.* Cerebral Cortex, 2001. **11**(7): p. 628-635.
304. Peterson, E.J. and C.A. Seger, *Many hats: intratrial and reward level-dependent BOLD activity in the striatum and premotor cortex.* Journal of Neurophysiology, 2013. **110**(7): p. 1689-1702.
305. Shigemune, Y., et al., *Remembering with gains and losses: effects of monetary reward and punishment on successful encoding activation of source memories.* Cerebral Cortex, 2014. **24**(5): p. 1319-1331.
306. Murty, V.P., K.S. Labar, and R.A. Adcock, *Threat of punishment motivates memory encoding via amygdala, not midbrain, interactions with the medial temporal lobe.* The Journal of Neuroscience, 2012. **32**(26): p. 8969-8976.
307. Murty, V.P., K.S. LaBar, and R.A. Adcock, *Distinct medial temporal networks encode surprise during motivation by reward versus punishment.* Neurobiol Learn Mem, 2016. **134 Pt A**: p. 55-64.
308. Seger, C.A., *The basal ganglia in human learning.* Neuroscientist, 2006. **12**(4): p. 285-90.
309. Seger, C.A. and C.M. Cincotta, *The roles of the caudate nucleus in human classification learning.* The Journal of Neuroscience, 2005. **25**(11): p. 2941-2951.
310. Mushiake, H., M. Inase, and J. Tanji, *Neuronal activity in the primate premotor, supplementary, and precentral motor cortex during visually guided and internally determined sequential movements.* J Neurophysiol, 1991. **66**(3): p. 705-18.
311. Hardwick, R.M., et al., *Multimodal connectivity of motor learning-related dorsal premotor cortex.* Neuroimage, 2015. **123**: p. 114-28.

312. Tomassini, V., et al., *Diffusion-weighted imaging tractography-based parcellation of the human lateral premotor cortex identifies dorsal and ventral subregions with anatomical and functional specializations.* J Neurosci, 2007. **27**(38): p. 10259-69.
313. Keele, S.W., et al., *The cognitive and neural architecture of sequence representation.* Psychol Rev, 2003. **110**(2): p. 316-39.
314. Schubotz, R.I. and D.Y. von Cramon, *Functional-anatomical concepts of human premotor cortex: evidence from fMRI and PET studies.* Neuroimage, 2003. **20**: p. S120-S131.
315. King, B.R., et al., *Cerebral Activation During Initial Motor Learning Forecasts Subsequent Sleep-Facilitated Memory Consolidation in Older Adults.* Cerebral Cortex, 2017. **27**(2): p. 1588-1601.
316. Breton, J. and E.M. Robertson, *Dual enhancement mechanisms for overnight motor memory consolidation.* Nature human behaviour, 2017. **1**(6).
317. Carbon, M., et al., *Caudate nucleus: influence of dopaminergic input on sequence learning and brain activation in Parkinsonism.* Neuroimage, 2004. **21**(4): p. 1497-507.
318. Albouy, G., et al., *Both the hippocampus and striatum are involved in consolidation of motor sequence memory.* Neuron, 2008. **58**(2): p. 261-272.
319. Debas, K., et al., *Off-line consolidation of motor sequence learning results in greater integration within a cortico-striatal functional network.* Neuroimage, 2014. **99**: p. 50-58.
320. Brown, R.M. and E.M. Robertson, *Off-line processing: reciprocal interactions between declarative and procedural memories.* J Neurosci, 2007. **27**(39): p. 10468-75.
321. Floyer-Lea, A. and P.M. Matthews, *Changing brain networks for visuomotor control with increased movement automaticity.* J Neurophysiol, 2004. **92**(4): p. 2405-12.
322. Kable, J.W. and A. Chatterjee, *Specificity of action representations in the lateral occipitotemporal cortex.* J Cogn Neurosci, 2006. **18**(9): p. 1498-517.
323. Kable, J.W., et al., *Conceptual representations of action in the lateral temporal cortex.* J Cogn Neurosci, 2005. **17**(12): p. 1855-70.
324. Lingnau, A. and P.E. Downing, *The lateral occipitotemporal cortex in action.* Trends Cogn Sci, 2015. **19**(5): p. 268-77.
325. Saga, Y., et al., *Ventral Pallidum Encodes Contextual Information and Controls Aversive Behaviors.* Cereb Cortex, 2017. **27**(4): p. 2528-2543.
326. Faget, L., et al., *Opponent control of behavioral reinforcement by inhibitory and excitatory projections from the ventral pallidum.* Nat Commun, 2018. **9**(1): p. 849.
327. Robertson, E.M., A. Pascual-Leone, and R.C. Miall, *Current concepts in procedural consolidation.* Nature Reviews. Neuroscience, 2004. **5**(7): p. 576-582.
328. Albouy, G., et al., *Daytime sleep enhances consolidation of the spatial but not motoric representation of motor sequence memory.* PLoS One, 2013. **8**(1): p. e52805.
329. Byrne, J.E.M., et al., *Time of day differences in neural reward functioning in healthy young men.* The Journal of Neuroscience, 2017. **37**(37): p. 8895-8900.
330. Steel, A., C. Thomas, and C.I. Baker, *Effect of time of day on reward circuitry. A discussion of Byrne et al. 2017.* arXiv preprint arXiv:1802.00753, 2018.
331. Kolasinski, J., et al., *Investigating the Stability of Fine-Grain Digit Somatotopy in Individual Human Participants.* J Neurosci, 2016. **36**(4): p. 1113-27.
332. Smith, S.M., et al., *Correspondence of the brain's functional architecture during activation and rest.* Proc Natl Acad Sci U S A, 2009. **106**(31): p. 13040-5.
333. van den Heuvel, M.P. and H.E. Hulshoff Pol, *Exploring the brain network: a review on resting-state fMRI functional connectivity.* European Neuropsychopharmacology, 2010. **20**(8): p. 519-534.

334. Edden, R.A., et al., *Orientation discrimination performance is predicted by GABA concentration and gamma oscillation frequency in human primary visual cortex.* J Neurosci, 2009. **29**(50): p. 15721-6.
335. Yoon, J.H., et al., *GABA concentration is reduced in visual cortex in schizophrenia and correlates with orientation-specific surround suppression.* J Neurosci, 2010. **30**(10): p. 3777-81.
336. Robertson, C.E., E.M. Ratai, and N. Kanwisher, *Reduced GABAergic Action in the Autistic Brain.* Curr Biol, 2016. **26**(1): p. 80-5.
337. Zhang, Y., M. Brady, and S. Smith, *Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm.* IEEE Transactions on Medical Imaging, 2001. **20**(1): p. 45-57.
338. Lauzon, M.L. and B.K. Rutt, *Polar sampling in k-space: reconstruction effects.* Magn Reson Med, 1998. **40**(5): p. 769-82.
339. Klose, U., *In vivo proton spectroscopy in presence of eddy currents.* Magn Reson Med, 1990. **14**(1): p. 26-30.
340. Walsh, D.O., A.F. Gmitro, and M.W. Marcellin, *Adaptive reconstruction of phased array MR imagery.* Magn Reson Med, 2000. **43**(5): p. 682-90.
341. Dreher, W. and D. Leibfritz, *New method for the simultaneous detection of metabolites and water in localized in vivo 1H nuclear magnetic resonance spectroscopy.* Magn Reson Med, 2005. **54**(1): p. 190-5.
342. Soher, B.J., et al., *Vespa: Integrated applications for RF pulse design, spectral simulation and MRS data analysis.* ISMRM, 2011: p. 1410.
343. Desikan, R.S., et al., *An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest.* Neuroimage, 2006. **31**(3): p. 968-80.
344. O'Shea, J., et al., *Functionally specific reorganization in human premotor cortex.* Neuron, 2007. **54**(3): p. 479-90.
345. Vul, E., et al., *Puzzlingly High Correlations in fMRI Studies of Emotion, Personality, and Social Cognition.* Perspectives on Psychological Science, 2009. **4**(3): p. 274-290.
346. Near, J., et al., *Unedited in vivo detection and quantification of γ-aminobutyric acid in the occipital cortex using short-TE MRS at 3 T.* NMR in Biomedicine, 2013. **26**(11): p. 1353-1362.
347. Duncan, N.W., C. Wiebking, and G. Northoff, *Associations of regional GABA and glutamate with intrinsic and extrinsic neural activity in humans—a review of multimodal imaging studies.* Neuroscience and Biobehavioral Reviews, 2014. **47**: p. 36-52.
348. Stagg, C.J., V. Bachtiar, and H. Johansen-Berg, *The role of GABA in human motor learning.* Curr Biol, 2011. **21**(6): p. 480-4.
349. Thielen, J.-W., et al., *The increase in medial prefrontal glutamate/glutamine concentration during memory encoding is associated with better memory performance and stronger functional connectivity in the human medial prefrontal-thalamus-hippocampus network.* Human Brain Mapping, 2018. **39**(6): p. 2381-2390.
350. Bezalel, V., R. Paz, and A. Tal, *Inhibitory and excitatory mechanisms in the human cingulate-cortex support reinforcement learning: A functional Proton Magnetic Resonance Spectroscopy study.* Neuroimage, 2018. **184**: p. 25-35.
351. Rubia, K., et al., *Methylphenidate normalises activation and functional connectivity deficits in attention and motivation networks in medication-naïve children with ADHD during a rewarded continuous performance task.* Neuropharmacology, 2009. **57**(7-8): p. 640-652.
352. Konrad, K. and S.B. Eickhoff, *Is the ADHD brain wired differently? A review on structural and functional connectivity in attention deficit hyperactivity disorder.* Human Brain Mapping, 2010. **31**(6): p. 904-916.

353. Boer, V.O., et al., *Multislice ^{1}H MRSI of the human brain at 7 T using dynamic B_0 and B_1 shimming.* Magnetic Resonance in Medicine, 2012. **68**(3): p. 662-670.
354. Goelman, G., et al., *Chemical-shift artifact reduction in Hadamard-encoded MR spectroscopic imaging at high (3T and 7T) magnetic fields.* Magnetic Resonance in Medicine, 2007. **58**(1): p. 167-173.
355. Hangel, G., et al., *Ultra-high resolution brain metabolite mapping at 7 T by short-TR Hadamard-encoded FID-MRSI.* Neuroimage, 2018. **168**: p. 199-210.
356. Reardon, P.K., et al., *Normative brain size variation and brain shape diversity in humans.* Science, 2018. **360**(6394): p. 1222-1227.
357. Seidlitz, J., et al., *Morphometric Similarity Networks Detect Microscale Cortical Organization and Predict Inter-Individual Cognitive Variation.* Neuron, 2018. **97**(1): p. 231-247.e7.
358. Westwater, M.L., et al., *Associations between cortical thickness, structural connectivity and severity of dimensional bulimia nervosa symptomatology.* Psychiatry Research: Neuroimaging, 2018. **271**: p. 118-125.
359. Smith, S.M., et al., *A positive-negative mode of population covariation links brain connectivity, demographics and behavior.* Nature Neuroscience, 2015. **18**(11): p. 1565-1567.
360. Murty, V.P., et al., *Selectivity in post-encoding connectivity with high-level visual cortex is associated with reward-motivated memory.* J Neurosci, 2016.
361. Murayama, Y., et al., *Relationship between neural and hemodynamic signals during spontaneous activity studied with temporal kernel CCA.* Magnetic Resonance Imaging, 2010. **28**(8): p. 1095-1103.
362. Owen, A.M., et al., *Putting brain training to the test.* Nature, 2010. **465**(7299): p. 775-8.
363. Haith, A.M., J. Pakpoor, and J.W. Krakauer, *Independence of Movement Preparation and Movement Initiation.* J Neurosci, 2016. **36**(10): p. 3007-15.
364. Garrett, D., S. MacDonald, and F. Craik, *Intraindividual reaction time variability is malleable: feedback- and education-related reductions in variability with age.* Frontiers in Human Neuroscience, 2012. **6**(101).
365. Wong, A.L., et al., *Reaction times can reflect habits rather than computations.* Elife, 2017. **6**.
366. Muhammed, K., et al., *Reward sensitivity deficits modulated by dopamine are associated with apathy in Parkinson's disease.* Brain, 2016. **139**(Pt 10): p. 2706-2721.
367. Reimer, J., et al., *Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex.* Nat Commun, 2016. **7**: p. 13289.
368. Karson, C.N., *Spontaneous eye-blink rates and dopaminergic systems.* Brain, 1983. **106 (Pt 3)**: p. 643-53.
369. Quattrocchi, G., et al., *Pharmacological Dopamine Manipulation Does Not Alter Reward-Based Improvements in Memory Retention during a Visuomotor Adaptation Task.* eNeuro, 2018. **5**(3).
370. Jackson, G.M., et al., *Serial reaction time learning and Parkinson's disease: Evidence for a procedural learning deficit.* Neuropsychologia, 1995. **33**(5): p. 577-593.
371. Semrau, J.A., J.S. Perlmutter, and K.A. Thoroughman, *Visuomotor adaptation in Parkinson's disease: effects of perturbation type and medication state.* J Neurophysiol, 2014. **111**(12): p. 2675-87.
372. Matthis, J.S., J.L. Yates, and M.M. Hayhoe, *Gaze and the Control of Foot Placement When Walking in Natural Terrain.* Current Biology, 2018. **28**(8): p. 1224-1233.e5.
373. Mathis, A., et al., *DeepLabCut: markerless pose estimation of user-defined body parts with deep learning.* Nat Neurosci, 2018. **21**(9): p. 1281-1289.
374. Wesseling, C., *Jets' Christian Hackenberg overhauls throwing motion*, in *Around the NFL*. 2018, National Football League: www.nfl.com.