Imperial College London

Deep Learning for Force Decoding from Multimodal Non-Invasive Brain Signals

Pablo Ortega San Miguel
2021

Supervised by
Prof. A Aldo Faisal

Department of Computing
Imperial College London

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Computer Science of Imperial College London
and the Diploma of Imperial College London

Declaration of Originality

I hereby confirm that the work presented in this thesis is my own, and that appropriate references have been used to denote the work of others. Whenever a piece of work has been carried out in collaboration with someone else their contribution is also specified.

Pablo Ortega San Miguel

January, 2021.

Copyright Statement

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Acknowledgements

I want to thank my supervisor, Aldo Faisal, for giving me this opportunity and guiding me in this arduous process. In your lab, I have found the authentic experience of mixing neuroscience and engineering that I always looked for. You have taught me to navigate the scientific jungle. Only now I start to see where I departed from, where I am at, what surrounds me, and where I want to go. This would not have been possible with the tools and the experience you provided me with. Especially, thanks for your support and your continuous guidance during pandemic times. It gave me the courage and faith to keep going in the midst of uncertainty.

I also want to thank Shlomi Haar. I will never forget the conferences that we assisted together and all our discussions about motor control. Thank you, Tong Zhao, for helping me collect data. You made those endless hours seem minutes, and I hope I could teach you as many things as I learned from you. Xiaoxi Wei, I immensely enjoyed working with you on our BCI demonstration for the Cybathlon. I am looking forward to seeing your work featuring in the next edition. It was a pleasure discussing new BCI directions with you and spending our BCI hack-days together. I am very grateful for everyone's help in the lab and making, day by day, everything easier. Harris Konnaris and Alex Harston, you have been there for me in the most challenging times, and your words always cheered me up.

I also want to express my gratitude to Krysia Broda, Amani El-Kholy, Alessandra Russo, Patel Hassan, Nicola Hann and everyone in the Department of Computing and the HiPEDS CDT. Thanks for making the CDT possible and for your support every day. Again, thank you for your remarkable work during the pandemic, helping my peers and me getting through it in the best way possible.

To all my Londoners friends, Lou Ann, Patri, Juan, Joan, Adrian, Mireia, Hugo, Quentin, Pieri, Hannah, Angu, Axel and many more I, you are the most loving, and you have made me not want to leave London. Especially, Tom Nabarro, you are a constant source of inspiration and admiration. Thank you for your collaboration at work and your friendship.

Thanks very much to my parents and all my family in Spain for always supporting me in my journey. Last and most important, thanks to my chosen family and *quaranteam* in London: Irene and Marina and my partner Katie. You have been there in literally all my ups and downs and have endured me closely since the beginning. That is a lot, and I would have never asked anyone for so much, but I did not need to ask you. Katie, I owe so many adventures. Thank you for being so loving, patient, understanding and endlessly supportive. There is not a day I do not feel lucky to have you.

Many people, from many nations, during uncountable days and in tremendously different ways, have enabled this research. I am looking forward to keeping working, teach and give at least the same support and opportunities to those coming tomorrow. Love and science are not contained by borders. In the current political scenario, ideas and people might. Let us not spare a moment to work together as one.

Dedication

Dedicated to Katie Goss, Irene García & Marina Saiz.
You have been consistently supportive throughout
and when I most needed.

'I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.'

Alan Turing, Computing Machinery and Intelligence (1950)



The Universal Turing Machine, by Jin Wicked

Abstract

Brain-Computer Interfaces (BCIs) bypass the nerve-muscle communication pathways. They offer people with disabilities new opportunities to assist, rehabilitate and restore their force control. Such control allows us to stabilise our manipulations, making them precise and safe. However, non-invasive BCIs, a surgery-free technology, have only made moderate progress in force decoding compared to invasive approaches. In this thesis, we explore and develop deep learning (DL) force decoding methods for electroencephalography (EEG) and functional Near-Infrared Spectroscopy (fNIRS): two signals reflecting complementary brain function information.

We first release HYGRiP, a dataset including ready-to-use multimodal brain (fNIRS and EEG) and unilateral hand signals (force and electromyography, EMG). We show that fast force generation introduces non-stationarities that remarkably increase fNIRS variability. Hence, we developed hemCNN, a convolutional neural network (CNN) method for fNIRS signals decoding and neuroimaging that outperforms conventional decoding approaches and reveals hand-specific cortical activations.

We continue generating HYGRiiP. Here, both hands contract simultaneously but generate different dynamic forces. With this multimodal dataset, we aim to answer two longstanding questions in non-invasive BCI: the possibility to decode dynamic forces and two simultaneous degrees of freedom (DoF). We show that multimodal non-invasive brain signals reconstruct bimanual force better than either of the signals alone. Our deep-fusion method further increases force reconstruction and better exploits signals with different physiological and physical origins. We show that most of the reconstruction is due to the robust detection of hand-specific contraction, which outperforms previous linear and EEG methods. Although we identified a limited hand-specific reconstruction of dynamic forces, our deep-fusion method revealed traces of hand-specific force modulation in EEG and fNIRS, suggesting that force modulated signals are accessible without surgery. However, current sensing systems are still too noisy to detect them on a single-trial basis.

Our results show that appropriately designed DL methods can improve fNIRS and multimodal signal processing for force decoding.

Contents

Acknowledgements **7**

Abstract **10**

1 Introduction **23**

1.1 Motivation and Objectives 23

1.2 Contributions 26

1.3 Thesis structure 28

2 Background **30**

2.1 Human motor control 30

2.2 Non-invasive and portable Brain-Computer Interfaces 34

2.3 Force decoding in BCIs: A systematic review 44

3 HYGRiP: Multimodal full brain-muscle characterisation of unimanual grip **68**

3.1 Introduction 68

3.2 Methods and materials 70

3.3 Data validation 79

3.4 Conclusion 85

4 HemCNN: Improving fNIRS' temporal & spatial resolution with Deep Learning **86**

4.1 Introduction 86

4.2 Methods 88

4.3 Results 101

4.4 Discussion 109

4.5 Conclusion 113

5 HYGRiiP: Multimodal full brain-muscle characterisation of bimanual grip 114

5.1 Introduction . . . 114

5.2 Methods and materials . . . 118

5.3 Data validation . . . 120

5.4 Conclusion . . . 125

6 Decoding bimanual force from multimodal brain signals with Deep Learning 126

6.1 Introduction . . . 126

6.2 Methods . . . 129

6.3 Results . . . 135

6.4 Discussion . . . 150

6.5 Conclusion . . . 154

7 Conclusion 158

7.1 Thesis achievements in context . . . 158

7.2 Immediate applications . . . 163

7.3 Future work . . . 165

Bibliography 169

List of Tables

2.1 Top five relevant fNIRS processing methods by number of citations 42

2.2 Comparison of queried studies specifically targeting force decoding 47

2.3 Force related information per brain signal 50

3.1 HYGRiP characteristics 69

4.1 Filters shapes and the total number of weights for hemCNN 95

4.2 HYGRiP benchmark 102

5.1 HYGRiiP and HYGRiP short comparative 116

6.1 Comparison of the EEG, fNIRS and multimodal systems performance in reconstructing force trajectories 136

6.2 Summary of bimanual force decoding methods comparison 140

6.3 Statistical measures for the binary detection of force (relaxation vs contraction). All deep models outperform the linear method in force detection (light grey) and have the highest performance (green). In particular, the *cnnatt-lstm* presents the highest F1 score and is the best choice for force detection purposes. 156

6.4 Comparison of correlation between real and decoded force trajectories from EEG, fNIRS and multimodal (fNIRS & EEG) systems using a linear model 156

6.5 Comparison of force reconstruction for the EEG, fNIRS and multimodal (fNIRS & EEG) systems using a linear model 157

6.6 Force decoding methods comparison 157

List of Figures

2.1 Motor control . . . 31
2.2 BCI schematic . . . 34
2.3 fNIRS instrumentation and Haemodynamic Response physiology . . . 37
2.4 The Haemodynamic Response . . . 39
2.5 Confounding sources of Haemoglobin changes . . . 40
2.6 Simulated non-stationary haemodynamic responses . . . 41
2.7 Accumulated publications in the BCI force decoding field from 2003 until October 2020 46

3.1 HYGRiP experimental setup . . . 71
3.2 Detail of the custom 3D holder design to co-align fNIRS and EEG sensors . . . 74
3.3 HYGRiP EEG & fNIRS data . . . 78
3.4 HYGRiP EMG data . . . 79
3.5 Validation of HYGRiP . . . 80
3.6 Multimodal signal periodogram for HYGRiP . . . 83

4.1 HemCNN architecture . . . 96
4.2 HemCNN leave-one-subject-out training . . . 97
4.3 Contraction and HR during the left/right hand-grip task . . . 100
4.4 Contraction evaluation and HR lateralisation during the left/right hand-grip task . . . 101
4.5 Leave-one-subject-out test accuracy across subjects for hand use classification . . . 104
4.6 HemCNN representative interpretation . . . 105
4.7 HemCNN cross-subject interpretation . . . 108

5.1 HYGRiiP experimental setup . . . 117
5.2 HYGRiiP EEG & fNIRS data . . . 121
5.3 HYGRiiP EMG . . . 122
5.4 Validation of HYGRiiP measures . . . 123

5.5 Scalp hemodynamics have negligible force-predicting power 124

6.1 Bimanual task, decoding features and models 132

6.2 Performance of bimanual force decoding for different metrics, methods and sensor modalities 136

6.3 Representative force trajectories reconstructed by different methods 139

6.4 Sensitivity analysis 142

6.5 Saliency time course 144

6.6 Topoplot of saliency associated to each hand 146

6.7 Topoplot of saliency associated to each generated force 148

6.8 Saliency during dynamic force generation 149

Acronyms

ALS Amyotrophic Lateral Sclerosis

ANN Artificial Neural Network

BCI Brain-Computer Interface

BMI Brain-Machine Interface

BTT Backpropagation Through Time

CC Correlation Coefficient

CF Convolutional Filter

CNS Central Nervous System

CNN Convolutional Neural Network

CW-fNIRS Continuous Wave functional Near-Infrared Spectroscopy

DBS Deep Brain Stimulation

DL Deep Learning

DoF Degrees of Freedom

dPM Dorsal Premotor Cortex

ECoG Electrocorticography

EEG Electroencephalography

ELU Exponential Linear Unit

EMG Electromiography

EOG Electrooculography

ERD Event Related Desynchronisation

ESN Echo State Networks

FDAcc Force Decoding Accuracy

fMRI functional Magnetic Resonance Imaging

fNIRS functional Near-Infrared Spectroscopy

FVAF Fraction of Variance Accounted For

GLM General Linear Model

GLM-hPCA General Linear Model with Hierarchical Principal Component Analysis

Hb Haemoglobin

HbO Oxygenated Haemoglobin

HbR Deoxygenated Haemoglobin

HDAcc Hand Disentanglement Accuracy

HDF Hierarchical-Data File

HR Haemodynamic Response

HRF Haemodynamic Response Function

HYGRiP Hybrid (fNIRS & EEG) unimanual (left & right hand) isometric Grip dataset

HYGRiiP Hybrid (fNIRS & EEG) bimanual isometric Grip dataset

ICA Independent Component Analysis

IPS Ipsilateral Intraparietal Sulcus

LDA Linear Discriminant Analysis

LFP Local Field Potential

LI Lateralisation Index

LMP Local Motor Potential

LSTM Long-Short Term Memory Network

M1 Primary Motor Cortex

MDN Mixture Density Network

ME Microelectrode

MLP Motor Local Potential

MSE Mean Squared Error

MVC Maximum Voluntary Contraction

NLP Natural Language Processing

PC Principal Component

PCA Principal Component Analysis

PMd Premotor dorsal cortex

PMv Premotor ventral cortex

PP Posterior parietal cortex

RLP Relevance Layer Propagation

RNN Recurrent Neural Network

S1 Primary Somatosensory Cortex

SCI Spinal Cord Injury

SCC Spearman's Correlation Coefficient

SD Source-Detector

SMA Supplementary Motor Area

SNR Signal-to-Noise Ratio

SPM Statistical Parametric Mapping

STN Subthalamic Nucleus

SVM Support Vector Machine

TD-fNIRS Time Domain functional Near-Infrared Spectroscopy

V1 Primary Visual Cortex

WHO World Health Organisation

Chapter 1

Introduction

1.1 Motivation and Objectives

We live in a physical world bounded by forces. Since we wake up till we go to sleep, we constantly interact with this world by applying and being subjected to forces. If you make your breakfast, even a simple coffee and toast would require you to operate the coffee machine, fill it with water, and use a knife to cut the bread. In both cases, your ability to generate adequate forces with both hands allowed you to feed yourself and do it without spilling coffee grain everywhere or cutting yourself. Or maybe not? Even now, you are probably generating forces while reading this. You may be taking notes by softly tapping on a keyboard or firmly holding a tablet with one hand while you precisely highlight with the other.

Forces define our daily and most basic interactions (Westling and Johansson, 1984; Ostry and Feldman, 2003). As humans, the specie with the most extraordinary dexterity, one of the main ways we interact with the world is our hands. From a physical point of view, forces are the cause of movements. From a biological perspective, we regulate the forces we generate to feed ourselves, adapt to the environment, and ultimately survive and reproduce. Psychologically, we may have different motivations and intentions, driving us to learn or improve a particular task requiring fine control of forces, for instance, playing a musical instrument or skating on ice. In any case, our brain enables us to learn to move around and survive in the world. Human motor control is the field that studies how we regulate our movement at the mechanical, physiological and psychological level.

However, not everyone can perform the same level of motor control (Canning et al., 1999; Anderson,

2004; Ang et al., 2015). The World Health Organisation (WHO) estimates that every year, around 250 to 500 thousand people suffer an Spinal Cord Injury (SCI) (WHO, 2002). In 2019 the Spinal Cord Injury Association (SIA) reported that half of the SCIs led to quadriplegia in the United Kingdom (Earl and Hartley, 2019). However, SCI is not the only reason leading to a motor disability. Amyotrophic Lateral Sclerosis (ALS) causes progressive degeneration of the peripheral nerves that gradually become worse at sending signals to the muscles. In particular, ALS is expected to increase from a 9% to a 16% incidence worldwide due to the increase of the population above age 60 (Arthur et al., 2016). Stroke is another condition whose incidence may increase due to the ageing of the population (Feigin et al., 2017). It is estimated that stroke causes motor and cognitive disability to 5 million people a year around the world (WHO, 2002).

In this context, people with motor disabilities can benefit from the use of Brain-Computer Interfaces (BCIs) (Daly and Wolpaw, 2008; Soekadar et al., 2011)[1]. Quoting Wolpaw et al. (2000),

> *"BCIs give their users communication and control channels that do not depend on the brain's normal output channels of peripheral nerves and muscles."*

The most remarkable advances in BCI have been made using invasive approaches (e.g. Bouton et al., 2016; Ganzer et al., 2020). Namely, a user undergoes surgery and sensors are chronically implanted inside the scalp. Using these sensors, we can record brain signals very close to their neuronal source (Carmena et al., 2003). Recently, an SCI patient was able to recover his hand control with this technology (Ganzer et al., 2020). While this is an outstanding demonstration, invasive BCIs are an expensive alternative that increases the burden put on the health of already health-compromised people due to the surgical procedures they require. Furthermore, not all disabilities are irreversible or require lifelong interventions (Soekadar et al., 2011; Silvoni et al., 2011; Kaiser et al., 2011). In 2011, the WHO estimated that an average person spends 8 years of their life with a disability, pointing out the directions to pursue to overcome the health problems associated with disability (WHO, 2011). The report remarked that we need more affordable assistive devices to overcome the health problems derived from disabilities.

Non-invasive BCIs offer a plug-and-play solution that does not pose a health risk (Wolpaw et al., 2000). Non-invasive BCI can be used at any time, can be easily worn, and users can take them off at

[1]In this thesis, we use the term Brain-Computer Interfaces (BCIs) instead of Brain-Machine Interfaces (BMIs). Ultimately, machines are controlled using microprocessors, by extension, computers. Thus, we consider that the term BCI also includes BMIs in a more general sense.

their will (Casson, 2019). Non-invasive approaches also come at a fraction of the cost of the invasive alternatives (Nicolas-Alonso and Gomez-Gil, 2012). Indeed, technical advances in material science and miniaturisation are making non-invasive BCI more portable, cheaper and closer to commodification (Casson, 2019; Kernel, 2020). Other technical advances are enabling the measurement of evermore sophisticated physical phenomena, opening new ways in which brain activity can be measured non-invasively (Torricelli et al., 2014; Lange and Tachtsidis, 2019; Lacerenza et al., 2020). For example, the first-ever non-invasive human brain activity recordings took place in 1924, measuring electrical potentials or Electroencephalography (EEG) with electrodes placed over the scalp (Haas, 2003). In the early 90s, magnetic resonance and near-infrared light were used to measure changes in haemoglobin concentration in the human brain leading to functional Magnetic Resonance Imaging (fMRI) (Kwong et al., 1992; Rao et al., 1993; Menon et al., 1995) and functional Near-Infrared Spectroscopy (fNIRS) (Hoshi and Tamura, 1993; Kato et al., 1993; Villringer et al., 1993). More recently, photonic measures are becoming more specialised in detecting neuronal activity (Phan et al., 2016; De Roever et al., 2017), and used in combination with acoustics (photoacoustics) to provide joint morphological and functional information (Xia et al., 2014; Tang et al., 2016).

Simultaneously, machine learning is day-by-day conquering tasks that only humans could accomplish a few years ago or that we thought would take years to be solved. In 2017, a machine won for the first time a human playing Go (Silver et al., 2017); in 2018, another algorithm learned to provide intensive care in a hospital setting (Komorowski et al., 2018); and in 2020, yet another one improved protein folding predictions (Senior et al., 2020), solving the problem years before expected. All these algorithms were built without needing to programme each detailed bit of the complex phenomena they process. In particular, they show the potential to harness the increasingly sophisticated physical phenomena used to non-invasively sense the complex brain processes that control our behaviour in health and disease.

In this context, non-invasive BCIs might soon become a reliable and widespread technology enabling people with disabilities to carry out daily activities efficiently and, ultimately, increasing their autonomy. However, there is still a lack of algorithms that automatise the processing and fusion of novel and well-established non-invasive brain-sensing modalities (Ahn and Jun, 2017; Hong et al., 2018). In particular, machine learning approaches lack interpretability, are mainly devised for classification and are primarily applied in the decoding of executed or imagined movements, paying less attention to force generation.

The main objective of this thesis is to advance in the understanding of how non-invasive brain signals

from different physical origins can be exploited using machine learning —Deep Learning (DL)(LeCun et al., 2015) in particular— to improve BCIs with force decoding applications. To this end, I use two well-established non-invasive brain signals showing the most significant potential for miniaturisation and wearability in the short term: EEG and fNIRS. In particular, before moving into the multimodal (EEG and fNIRS) analysis, I first show that fNIRS —which measures slow brain haemodynamics— is a suitable measure to analyse brain activity in strenuous motor control scenarios. Secondly, I aim to clarify if multimodal non-invasive brain signals provide an advantage at decoding continuous force signals. My third objective consists of developing DL algorithms that improve conventional methods in analysing and fusing fNIRS and multimodal brain signals. Last, by releasing and making my data readily available, I aim to stimulate and invite researchers in neuroscience to the field of machine learning and machine learning experts to the field of neuroscience and, ultimately, introduce them to the exciting and meaningful problem of machine learning applied to BCIs.

The machine learning methods presented here will enable the processing and exploitation of the current and the next generation of multimodal brain-sensing techniques. The marrying of sophisticated physical phenomena to sense brain activity with powerful data-driven approaches show the potential to unlock non-invasive BCI, making it more reliable for assistance in every-day tasks. With this research work, I aim to help people with disabilities regaining autonomy via more affordable and surgery-free non-invasive BCIs and enable better behaviourally driven neuroimaging methods for non-invasive brain-sensing techniques.

1.2 Contributions

To advance the field of non-invasive BCI,

- I review and analyse the field of fNIRS processing methods for human motor control.
- I identify a lack of research in fNIRS processing methods able to adapt to strenuous human force generation conditions.
- I collect and publish a ready-to-use dataset including multimodal (EEG and fNIRS) human brain signals during strenuous unimanual force generation.
- I develop hemCNN, a Convolutional Neural Network (CNN) method for the decoding and analysis of fNIRS during strenuous human hand force generation.

- I review and analyse the field of BCI to decode force.
- While it has been shown that multimodal (EEG and fNIRS) carry complementary information; I identify a lack of multimodal BCI approaches for the continuous decoding of simultaneous bimanual forces.
- We collect and publish a ready-to-use dataset including multimodal (EEG and fNIRS) human brain signals during dynamic bimanual force generation.
- I show that multimodal non-invasive brain signals offer better decoding of continuous bimanual force than each modality.
- I develop deep-fusion methods for decoding continuous bimanual forces and show that my *cnnatt*, a CNN architecture with residual and attention modules, improves the fusion of brain signals with different physical origins.
- Using *cnnatt*, I show that there are traces of hand-specific force modulation in the non-invasive brain signals used here.

These contributions have been or are in the process of being published in high impact journals and conferences:

- Ortega, P. & Faisal, A. A. Deep Learning multimodal fNIRS & EEG signals for bimanual grip force decoding. Submitted to the *Journal of Neural Engineering*.
- Ortega, P. & Faisal, A. A. HemCNN: Deep Learning enables decoding of fNIRS cortical signals in hand grip motor tasks. Submitted to *Neuroimage*.
- Ortega, P., Zhao, T., & Faisal, A. A. HYGRiiP: Full-stack characterisation of neurobehavioural signals during a bimanual grip force tracking task. In review with *Nature Scientific Data*.
- Ortega, P. & Faisal, A. A. (2021). Deep real-time decoding of bimanual grip force from multimodal EEG & fNIRS sensor fusion. In *2021 10th International IEEE/EMBS Conference on Neural Engineering*. IEEE.
- Ortega, P. & Faisal, A. A. (2021). HemCNN: Deep Learning enables decoding of fNIRS cortical signals in hand grip motor tasks. In *2021 10th International IEEE/EMBS Conference on Neural Engineering*. IEEE.

- Ortega, P., Zhao, T., & Faisal, A. A. (2020). HYGRIP: Full-stack characterisation of neurobehavioural signals (fNIRS, EEG, EMG, force and breathing) during a bimanual grip force control task. *Frontiers in Neuroscience*, 14:919.

- Ortega, P. & Faisal A. A. (2019). Deep learning time-resolves fNIRS cortical activation asymmetries during a force tracking task evoking non-stationary haemodynamic responses. In *2019 49th International Conference of the Society for Neuroscience*, program No. 760.15. SfN.

- Ortega, P. & Faisal A. A. (2019). Optimal identification of the haemodynamic response under non-stationary conditions. In *2019 UK Sensorimotor Conference*. UCL.

- Ortega, P. & Faisal A. A. (2018). Mutually informed dynamical models for fNIRS denoising based on EMG and EEG. In *2018 48th International Conference of the Society for Neuroscience*, program No. 225.16. SfN.

- Ortega, P., Zhao T. & Faisal A. A. (2018). Directed information relationship between EMG, EEG and fNIRS during isometric hand contractions. In *2018, 28th Annual Conference of the Society for the Neural Control of Movement*. NCM.

1.3 Thesis structure

In this chapter, I have motivated and contextualised the research carried out in this thesis. Specifically, we will be focusing on BCI decoders of force using non-invasive brain signals.

In Chapter 2, I introduce motor control and EEG and fNIRS non-invasive brain signals. I specifically focus on fNIRS and the conventional way fNIRS signals are processed. I end with a review of the current state of BCI for force decoding spanning both non-invasive and invasive studies in humans and animals. In this review, I identify 4 venues that lack exploration in the context of non-invasive BCI for force decoding: (1) Data-driven non-linear signal processing methods for fNIRS; (2) multimodal (fNIRS and EEG methods for continuous force decoding; (3) non-linear approaches for non-invasive force decoding; and (4) bimanual force decoding.

In Chapter 3, I present HYGRiP, a validated and ready-to-use multimodal dataset including fNIRS and EEG brain signals, and force and Electromiography (EMG) as behavioural signals during unimanual strenuous force generation. This task has previously challenged conventional neuroimaging methods. I

release these data to push the development of methods that are able to deal with the non-stationarities the generation of strenuous forces introduces. I believe HYGRiP helps to draw more machine learning experts and neuroscientists into the exciting problem of BCI.

In Chapter 4, I develop hemCNN, an interpretable DL data-driven method for fNIRS processing in motor control tasks that evoke non-stationary haemodynamic cortical responses. I demonstrate how hemCNN can deliver better decoding and interpretable neuroimaging results where conventional methods, using a fixed haemodynamic response model, fail.

In Chapter 5, I present HYGRiiP, a validated and readily exploitable dataset including multimodal non-invasive brain signals (fNIRS and EEG) collected and designed for the exciting problem of bimanual force control. The motor task in HYGRiiP required participants to generate hand-specific dynamic forces with both hands simultaneously. In particular, HYGRiiP introduces the challenge of 2 simultaneous degrees of freedom control of force in BCI.

In Chapter 6, I explore the fusion of multimodal brain signals (EEG and fNIRS) for the decoding of simultaneous bimanual forces. First, I see whether multimodal signals offer any advantage in decoding hand-specific forces during simultaneous bimanual force generation. Then, I develop DL methods aiming to improve the fusion of brain signals (EEG and fNIRS) with very different physical and physiological origins. Finally, I analyse what these data-driven methods learn from multimodal signals.

I end with Chapter 7, in which I present concluding remarks about my findings and suggest future research directions.

Chapter 2

Background

2.1 Human motor control

From meeting our most basic physiological needs to delivering the most beautiful artistic performance or showing affection, our actions involve mechanical interactions between our body and what surrounds us (Fig. 2.1A) (Rosenbaum, 2009). The human motor system makes these interactions possible by providing mechanical support, mechanical power and mechanisms to control them. Respectively, the motor system can be divided into the musculoskeletal system and a highly distributed nervous system.

On the one hand, the musculoskeletal system is an articulated system with many degrees of freedom that provides the body with high mobility, mechanical strength and protection for our vital organs. The skeleton provides mechanical support while the muscles, by linking different skeletal points and contracting, generate the torques generating the movement (bottom of Fig. 2.1B). The muscles are made of muscular fibres with different contraction properties (e.g. velocity and fatigue resistance). Bundles of varying numbers of muscle fibres are innervated by the same motor neuron configuring a motor unit, the basic unit of motor control.

When the motor neuron fires, the fibres contract, shortening the muscle's length, approximating the points to which the muscles are attached and producing the movement. The characteristics and dynamics of these contractions and how the motor neurons regulate them are complex and involve firing frequencies, velocity, fibre type and length, and fatigue, among others (Hill, 1938; Gordon et al., 1966; Burke et al., 1973; Scott et al., 1996; Cheng et al., 2000). On the other hand, the Central Nervous System (CNS) is in charge of coordinating these fibre muscle contractions to produce meaningful

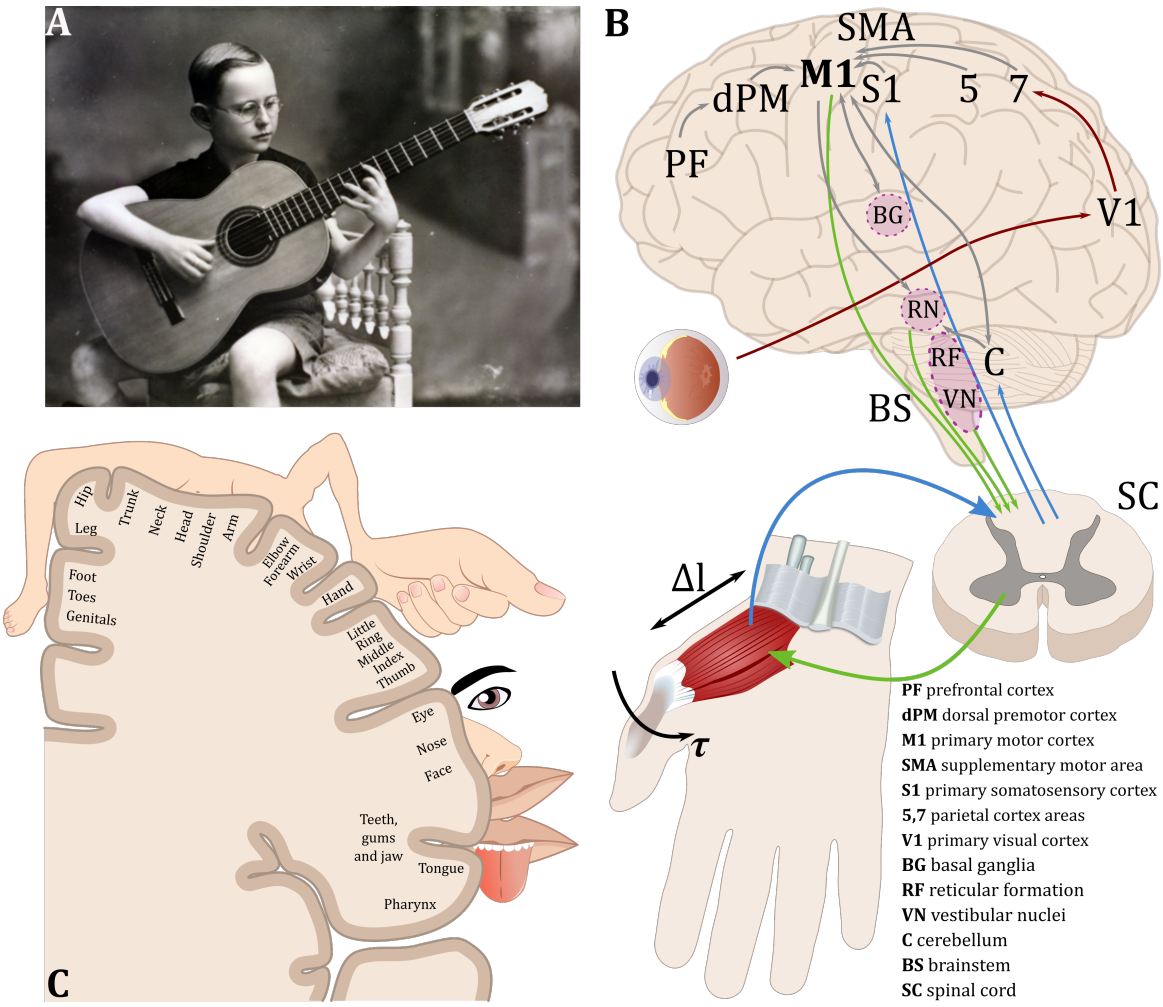


Figure 2.1: **A** A young Narciso Yepes plays the guitar (public domain from narcisoyepes.org). Playing the guitar is an example of sophisticated motor control involving the integration of visual, auditory, haptic and proprioceptive feedback to produce complex sound sequences with hand coordinated actions on the guitar. **B** Schematic of the musculoskeletal and neural elements involved in motor control (composition from modified Creative Commons elements: medulla by Pollarlis, hand by Wilfredo R, brain by Jon Philips and eye by Ryan Lerch). The muscle needs to contract (Δl) to generate the needed torque (τ) to pick the string with the right-hand thumb. This command is received in the form of firing rates through different motor units with synaptic connections with the Primary Motor Cortex (M1) in the spinal cord (SC) (green arrows). Inputs into M1 support the elaboration of the precise command. The cerebellum (C) has an important role in fine sequence executions like those required to play the guitar. Visual information (red arrows) is processed in the primary visual cortex (V1) and integrated in parietal areas (5,7) before being integrated in M1. Proprioceptive feedback (blue arrows) supplements the visual input and is processed in the primary somatosensory cortex (S1) before being integrated in M1. If Narciso has rehearsed several times, his prefrontal cortex (PF) might have some abstract representation of the piece that is translated into motor commands in the dorsal premotor cortex (dPM) and supplementary motor area (SMA). If he masters the piece, it is probably because the cerebellum has overtaken the PF creating a memory of precisely timed and coordinated finger executions. **C** Cortical motor homunculus as proposed by Penfield and Boldrey (1937) (Creative Commons by Popadius). The homunculus represents the somatotopic representation of the body in the motor cortex. In humans, with the highest hand dexterity across all species, the hand's cortical representation predominates compared with other body parts.

movements.

The CNS can be organised in three levels relative to motor control, with various degrees of conscious control characterising them (for a review, see Scott, 2016). From closer to the muscle and less conscious control (low level or primary control) to more conscious control (high level), these levels would be the spinal cord, the brainstem and the cerebral cortex (Figs. 2.1B and C). The spinal cord includes motor neurons and features low-level integration of ascending sensory information from joints and muscles. In particular, the spinal cord enables involuntary motor reflexes and schedules some fixed locomotion patterns. The brainstem provides further involuntary control of posture and also regulates the speed and quality of some locomotion patterns along with the cerebellum. The cerebral cortex provides the highest voluntary control level with a large and adaptable motor behaviour spectrum.

The motor behaviour can, at least, be divided into the planning and execution of motor actions. Different cortical areas have been associated with different roles spanning different degrees of movement control and preparation. M1 is primarily involved in goal-oriented motor execution and skilled motor tasks and projects the most significant number of axons to the corticospinal tract (Fig. 2.1B green). Most of these axonal projections regulate muscular contraction indirectly by modulating interneuron activities in the spinal cord. Interestingly, some of the M1 axonal projections directly regulate motor neuron activity through monosynaptic connections specially prevalent in hand muscles (Phillips, 1977; Porter and Lemon, 1993; McKiernan et al., 2000) and increasing in number with the level of dexterity across species (Heffner and Masterton, 1975; Nakajima et al., 2000).

This topographic and anatomical organisation is well captured by the homunculus (Fig. 2.1C) (Penfield and Boldrey, 1937). M1 also receives inputs from other brain areas providing somatosensory information, Primary Somatosensory Cortex (S1), and additional multimodal sensory information from the Primary Visual Cortex (V1), the Dorsal Premotor Cortex (dPM) or planning information from the prefrontal cortex, which are integrated in M1 and the premotor cortex (Scott, 2016). In particular, the Supplementary Motor Area (SMA) is very relevant for motor preparation and planning (Ball et al., 1999), being active even 1 s before the motor activity is executed and also when it is only imagined (Roland et al., 1980).

While the anatomical organisation and connectivity of M1 are well established (Penfield and Boldrey, 1937; Scott, 2016), the functional organisation at the neuronal level and its activity correlations with motor variables are less well understood and remains debated (Georgopoulos and Ashe, 2000; Todorov,

2000b). There is abundant evidence of neurons representing different hand trajectories direction and levels of force (Evarts, 2011). However, it has also been shown that the same neuronal activities of M1 correlate similarly to hand motions and other kinetic features in a non-specific manner (Mussa-Ivaldi, 1988; Todorov, 2000a; Sanger, 2003). In particular, this might result from the significant correlation in the motor variables (kinetic and kinematic) that define our movements.

Altogether, the motor system is a complex system encompassing a mechanical infrastructure and sophisticated ways of processing and integrating information for its control. Like motor neuron disease or spinal cord injury, several conditions disrupt the natural nerve-muscle pathway by which the brain controls our interactions with the environment. For many people suffering from motor disabilities, BCIs can provide a mean to regain control of their surroundings (Wolpaw et al., 2000; Daly and Wolpaw, 2008). These BCI systems can bypass the nerve-muscle communication pathways to enable the interaction with the external world through a computer or device (Wolpaw et al., 2000). As we have seen, the goal of the study of motor control is to understand the functional organisation at the neuronal level and model the role of different brain areas in the planning and executing movements. In contrast, from a BCI perspective, it is interesting to investigate the information originating in M1 that can be exploited for communication purposes.

While M1 is not necessarily the only source of information that users can use to control a BCI, it offers three advantages compared to exploiting other brain areas. First, compared to the spinal cord, the brainstem or the cerebellum, the information in M1 can be quickly and non-invasively accessed using affordable neuroimaging techniques such as EEG (e.g. Wolpaw et al., 2002; Soekadar et al., 2011; Schirrmeister et al., 2017) and fNIRS (e.g. Abdelnour and Huppert, 2009; Leff et al., 2011; Wriessnegger et al., 2017). Second, compared to other cortical areas like V1 that require the selective observation of external visual stimuli to generate a signal, M1, as we have seen, can be autonomously activated (Pfurtscheller and Lopes Da Silva, 1999; Pfurtscheller et al., 2006; Soekadar et al., 2011). Third, because many motor disabilities result from a malfunction of the lower parts of the corticospinal tract, the motor cortex remains functional but unused. Therefore it can be dedicated to BCI related control (Makin et al., 2015). This situation is especially relevant when the signals exploited are those naturally occurring in the cortex. In this way, the user does not have to learn new ways of generating these signals allowing a more intuitive control (Flint et al., 2014).

We here explore how to use non-invasive brain signals from the sensory-motor cortex to decode motor control features related to force. As we will see in the following sections, decoding force from non-

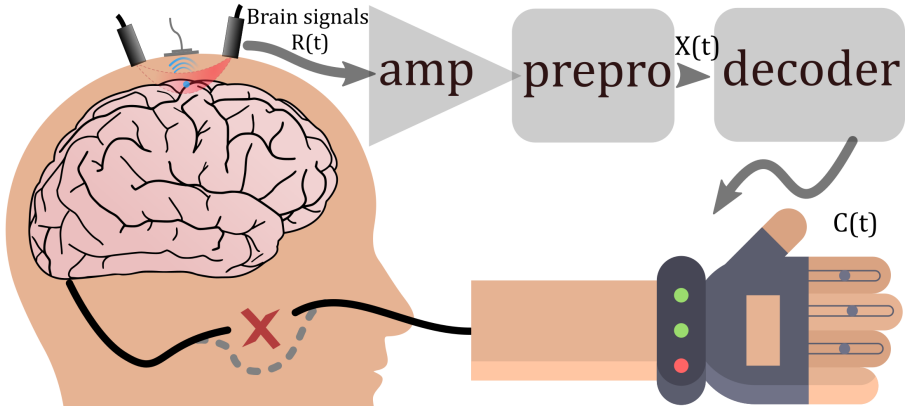


Figure 2.2: Non-invasive BCI scheme. Brain signals ($R(t)$, EEG and fNIRS in the figure) are non-invasively recorded. They are amplified and preprocessed ($X(t)$) and then fed to a decoding algorithm. The preprocessing may consist of filtering, normalisation, artefact removal or rejection and signal feature extraction. The decoding algorithm translates the preprocessed brain signals ($X(t)$) into commands that are sent to a machine (an assistive robotic hand in the figure). The control of that machine may serve different purposes, from completely recovering a lost function (e.g. hand-grip) to assisting in weakened motor functions (e.g. loss of grip force).

invasive brain signals poses some challenges. This work aims to advance force decoding in the BCI and neuroimaging directions.

2.2 Non-invasive and portable Brain-Computer Interfaces

As we have seen, motor control involves two complex systems, the musculoskeletal and nervous systems. These systems are integrated to regulate the movements we generate. Sometimes, the nervous system can fail to send commands to the muscles (Wolpaw et al., 2000). In this situation, BCIs –systems that translate brain signals into computer commands– can be used to control machines bypassing the natural communication pathway, i.e. nerves and muscles (Wolpaw et al., 2000). These machines can then be used to assist (Onose et al., 2012; Petković et al., 2012; Rayegani et al., 2014), rehabilitate (Daly and Wolpaw, 2008; Soekadar et al., 2011), or completely replace some motor function aspects (Bouton et al., 2016; Ganzer et al., 2020).

Depending on the instrumentation used to record the brain signals, BCIs can be classified as invasive and non-invasive. Invasive BCIs require surgical procedures to implant sensors inside the skull and include Electrocorticography (ECoG) meshes over the cortex or pia mater (Pistohl et al., 2013; Flint et al., 2014), UTAH arrays piercing the cortex (Ganzer et al., 2020) and Deep Brain Stimulation (DBS) sensors implanted in different brain areas, deep underneath the cortex (Khawaldeh et al., 2020). These systems have the obvious drawback of requiring expensive and complex surgeries (Wolpaw et al., 2000; McGie et al., 2013; Klein, 2016; Burwell et al., 2017). Due to biocompatibility issues, these invasive

systems also lead to the rejection of the implants that cause scarring around the sensors and the signal's deterioration with time (Klein, 2016). In contrast, non-invasive BCIs use recording systems that capture brain signals from the scalp without surgical procedures. Despite their lower Signal-to-Noise Ratio (SNR) (Ball et al., 2009; Buzsáki et al., 2012), non-invasive BCIs can be easily put on and taken off, do not have any health problems associated, are more affordable and, altogether, more suitable for commodification than invasive approaches (Nicolas-Alonso and Gomez-Gil, 2012; Casson, 2019).

Figure 2.2 depicts the elements of a non-invasive BCI. In the figure, brain signals, in particular fNIRS and EEG, are acquired non-invasively from sensors on the scalp. These raw signals are amplified and preprocessed to deliver more stable brain function features $X(t)$. Finally, a decoding algorithm translates these features or preprocessed signals $X(t)$ into a device's control commands $C(t)$. In the example of the figure, the device consists of a robotic hand which can be used to assist or rehabilitate hand movements when the nerve-muscle communication pathway deteriorates (discontinuous grey line) (Soekadar et al., 2015; Ang et al., 2015) or to completely replace muscular function when this pathway is fully broken (red cross) (Onose et al., 2012; Meng et al., 2016).

In particular, a BCI decoder is required to be causal, i.e. decoder only uses past and current information to decode information in the present or make future predictions. In the BCI terminology, these decoders can be tested in three ways: (1) offline, i.e. a portion of the recorded data not used to train the decoder, is used to test it; (2) online in open-loop, a trained decoder is used in real-time and translates brain signals into actions in a computer or the natural world through a machine while these actions are still performed in a natural way; and (3) online in closed-loop. This modality also requires the user to interact with the computer or the machine through brain signals.

The difference between the open and closed-loop cases during online control lies in how the feedback is delivered to the user, and it is easier to understand with an example. Imagine we ask a user to move their left or right hand while we record their brain signals. We then train an algorithm to decode the hand they intend to move from the brain signals. In the open-loop case, the users are still allowed to move their right or left hand, and the decoder only uses brain information to decode the hand. The users receive feedback on which hand they are moving via the BCI but also in a natural proprioceptive way since they can still feel their hands moving. In the closed-loop case, hand movements are not allowed or are not physically possible due to an actual or simulated lesion. Since there is no actual movement, the users only know what hand they are moving through the feedback delivered by the

BCI, hence closing the loop.

As we will see, most non-invasive BCI focus on the decoding of movements, but not force. However, force is essential in stabilising dynamical control during mechanical interactions (Franklin et al., 2003). For example, an action as simple as filling a glass of water would require different forces depending on the amount of water there is still in the jar and how we balance it. It has also been shown that the direct control of force leads to more naturalistic and robust BCI control in changing dynamic conditions (Cherian et al., 2011; Chhatbar and Francis, 2013). Indeed, if deteriorated, force should be rehabilitated as early as possible for successful motor recovery in the long-term after stroke (Canning et al., 1999; Cauraugh et al., 2000; Ada et al., 2006).

The work in this thesis is about offline BCI algorithms that decode the force that humans generate with their hands from non-invasive fNIRS and EEG signals. EEG and fNIRS are two non-invasive methods to measure brain activity. In comparison to other non-invasive methods like fMRI, fNIRS and EEG require more affordable instrumentation and setups that allow almost complete freedom of movement. In the following two sections, we provide enough background on these two techniques to understand how brain signals can be captured non-invasively with them.

2.2.1 Functional Near-Infrared Spectroscopy

Functional Near-Infrared Spectroscopy (fNIRS) is a non-invasive and portable neuroimaging technique that measures cortical brain activity by detecting the changes in Haemoglobin (Hb) concentration that follow neuronal activity (Cope et al., 1988; Leff et al., 2011; Ferrari and Quaresima, 2012). This spectroscopy technique was used with human neuroimaging purposes later than EEG and has increased in popularity since (Ferrari and Quaresima, 2012).

We briefly explain the working mechanism of the Continuous Wave functional Near-Infrared Spectroscopy (CW-fNIRS) system used here for data collection in this thesis and how it is related to brain activity. In particular, CW-fNIRS systems are smaller and more portable, facilitating their setup and the free movement of the subjects under study (Ferrari and Quaresima, 2012). Therefore, these systems are more appropriate for the study of motor function, BCI applications and commodity devices, compared to the constraining and more expensive fMRI scanners.

In particular, fNIRS offers a greater spatial resolution ($\leqslant 1\,\text{cm}$) than EEG (Ferrari and Quaresima, 2012) but has a lower temporal resolution, in the order of seconds, compared to EEG (milliseconds) (Buzsáki

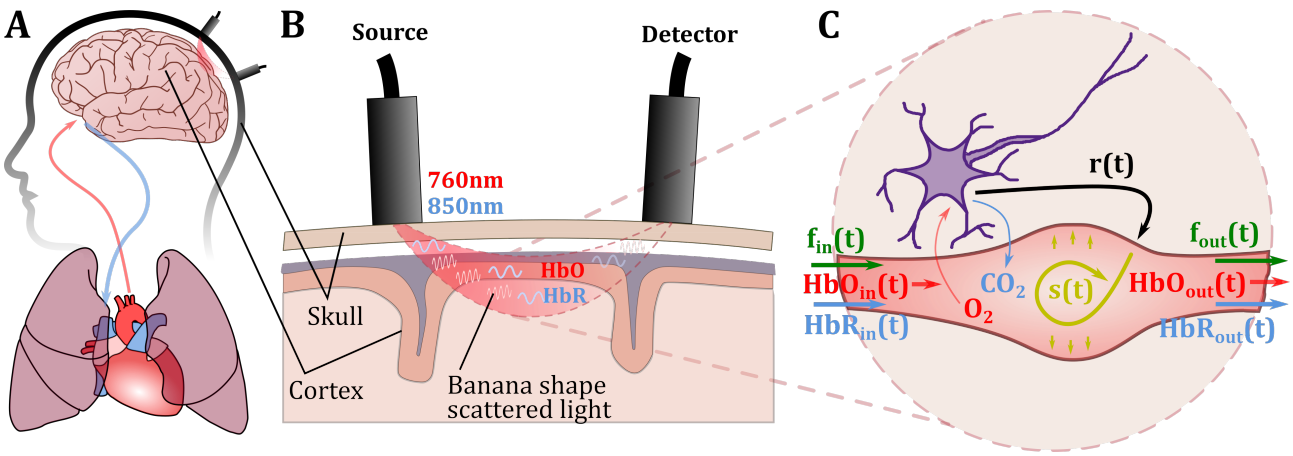


Figure 2.3: Schematic of the fNIRS instrumentation. **A** A source and a detector of fNIRS are placed over the scalp resulting in a banana-shaped infrared light-path that reaches the cortex where changes in Hb following neuronal activation can be measured. The lungs and heart also produce variations in blood oxygen levels in the brain and scalp. **B** Detail of the banana shape light path through the cortex. **C** Neurocoupling and Haemodynamic Response (HR). An active neuron demands oxygen to produce more energy. The depletion of oxygen levels in the surrounding bloodstream triggers a physiological response $r(t)$ that sends a vasodilatory signal to the capillaries $s(t)$ causing their dilation, thus increasing the blood flow ($f_{in}(t)$ and $f_{out}(t)$). A higher proportion of oxygen in form of Oxygenated Haemoglobin (HbO) overcompensate the previous oxygen depletion and is observed as an increase of HbO in fNIRS (see Fig. 2.4).

et al., 2012). Namely, while in EEG, the electrical fields' volume conduction makes it challenging to find the activity source (van den Broek et al., 1998; Leahy et al., 1998), in fNIRS it is guaranteed that the recorded activity is placed in between a sensor and a detector (Leff et al., 2011; Ferrari and Quaresima, 2012).

Light scattering and the Modified Beer-Lambert Law

Figure 2.3 shows a schematic of the fNIRS instrumentation used in this work. Low power light sources (5 – 10 mW) and detectors are placed on the scalp surface. Light scatters forming a banana shaped path (Okada and Delpy, 2003; Mansouri et al., 2010) from the source to the detector, reaching the cortical brain layers (Okada and Delpy, 2003; Custo et al., 2006) (Fig. 2.3A). The arc's depth in the banana-shaped scattering path is determined by the inter-optode (source-detector) distance (Okada and Delpy, 2003). Greater distances lead to deeper arcs (and deeper tissue layers) but also lower SNR since most photons are absorbed in the way and do not reach the detector above thermal noise levels (Mansouri et al., 2010) (Fig. 2.3B). Approximately, 3 cm inter-optode separation renders a penetration depth of 1.5 – 2 cm reaching 5 – 10 mm of cortical brain tissue where important cognitive processes take place (Okada and Delpy, 2003; Custo et al., 2006).

Along their travelling path, the scattered photons are selectively absorbed by particular molecules depending on their wavelength (Cope et al., 1988; Ferrari and Quaresima, 2012). Especifically, fNIRS

targets concentrations of Haemoglobin (Hb), the oxygen carrier in the blood. The Hb in the blood has two states: oxygenated (HbO) and deoxygenated or reduced Deoxygenated Haemoglobin (HbR), depending on whether it is transporting oxygen or not, respectively. In particular, each Hb specie preferentially absorbs photons of a specific wavelength (Wray et al., 1988) (Fig. 2.3B). By using two wavelengths of light, 760 and 850 nm, we can resolve changes in concentration of the two Hb types, since HbO absorbs more photons with 850 nm wavelength and less with 760 nm.

The Modified Beer-Lambert law (Delpy et al., 1988) is used to analytically find the amount of light that is absorbed at each wavelength and compute Hb changes of concentration.

$$\hat{I}_{ij}^{\lambda}(t) = \int_{\Gamma} \hat{I}_{ij}^{\lambda}(0) e^{-\mu_{abs}(r,t)} dr \tag{2.1}$$

where Γ is the path travelled by the light between the source i and the detector j; r is the coordinates vector; μ_{abs} is the coefficient of absorption of the tissue at each coordinate, $\hat{I}_{ij}^{\lambda}(t)$ is the estimated light intensity at the detector and $\hat{I}_{ij}^{\lambda}(0)$ the light intensity at the source at time $t = 0$. Namely, the equation reflects that the number of photons (light intensity) decays exponentially according to a μ rate (coefficient of absorption) as the light travels through a path Γ. The composition of the media ($\mu_{abs}(r,t)$) reflects the different light-absorbing tissues involved, such as the skin, bones, fat and, particularly, the blood vessels carrying Hb.

As mentioned, the use of two different wavelengths (λ), HbO absorbing more at 850 nm and less at 760 nm, makes it possible to separate the amount of light absorbed by each Hb type. The light intensity is normalised to an initial value to partially correct for nonlinearities and renders a measure called *differential optical density* (ΔOD),

$$\Delta OD_{ij}^{\lambda}(t) = Ln\left(\frac{\hat{I}_{ij}^{\lambda}(0)}{\hat{I}_{ij}^{\lambda}(t)}\right) \tag{2.2}$$

In practice, the lack of knowledge about the exact composition of the scattering medium, i.e. $\mu_{abs}(r)$, makes it necessary to make generalised assumptions (Ye et al., 2009). Population studies on the medium composition give ranges of values for $\mu_{abs}(r)$ and other coefficients (Duncan et al., 1995; Matcher et al., 1995) which, at the same time, allow to linearise Equation 2.2 for each wavelength λ,

$$\Delta OD_{ij}^{\lambda} = L_{ij}^{\lambda} DPF^{\lambda}(\epsilon_{\text{HbR}}^{\lambda}\Delta[\text{HbR}] + \epsilon_{\text{HbO}}^{\lambda}\Delta[\text{HbO}]). \tag{2.3}$$

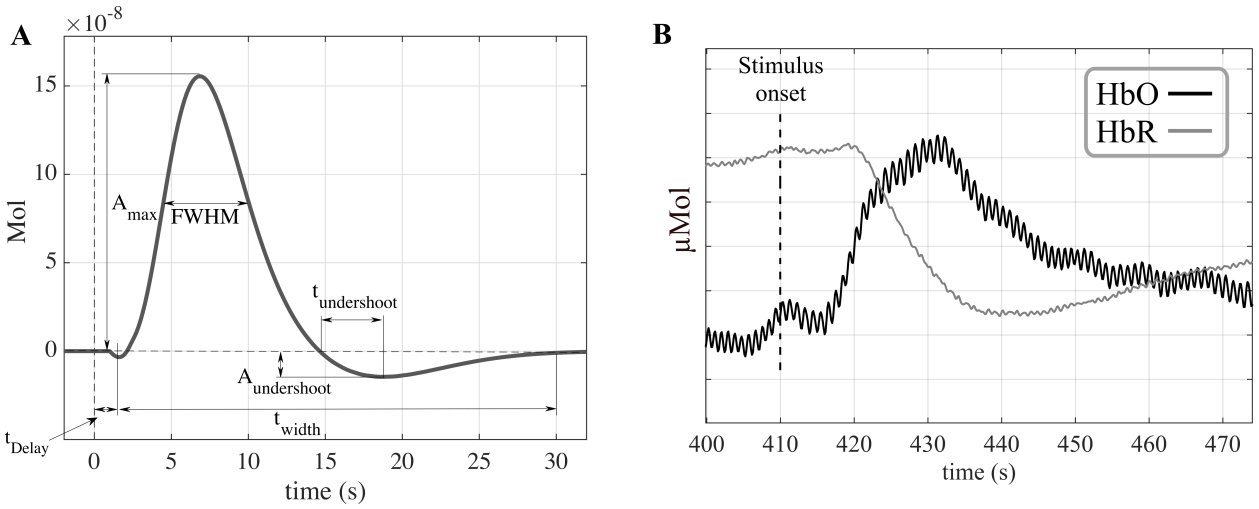


Figure 2.4: **A** Parameters defining the canonical HR for HbO, such as onset delay (t_{Delay}), amplitude (A_{max}) and undershoot of response, modelled using the sum of two gamma functions (Schroeter et al., 2004; Plichta et al., 2006; Ye et al., 2009) in molar units (M). FWHM, full-width half-maximum, i.e. width of the response at half its amplitude. **B** Real HR with HbO (black) and HbR (light grey). Not scaled units. We can observe the drift between the measures as different baseline levels before the stimulation. This drift is typically corrected during preprocessing but requires knowing when the stimulus or task start. The real HR is less clear than the model on the left. Many methods, like SPM, regress this real HR against models similar to those in the left to remove sources of noise. However, these models require selecting the parameters indicated in the figure, making these approaches very sensitive to the experimenter's subjective decisions (see main text).

In Equation 2.3: ΔOD_{ij}^{λ} is the change in optical density referred to the baseline at wavelength λ in cm^{-1}; L_{ij}^{λ} is the inter-optode distance per wavelength; DPF^{λ} is the dimensionless *differential path-length factor* accounting for the reduction in intensity due to scattering tissues; ϵ_{HbX}^{λ} is the *molar extinction coefficient* for each Hb specie in $M^{-1}cm^{-1}$ accounting for the light absorbed by each Hb type; and [HbX] is the molar concentration of each Hb type, X being either O for oxygenated, or R for reduced. From Equation 2.3, it is now clear that the use of two different wavelengths permits a solution for the linear system with two unknowns, the concentration of the two Hb types.

Neurocoupling and the Haemodynamic Response

We have seen how it is possible to measure changes in Hb concentration using fNIRS but, how are these Hb concentration changes related to neural activity? Figure 2.3C depicts how the activity of a neuron triggers a HR leading to changes in HbO and HbR concentration (Buxton et al., 1998; Friston et al., 2000; Buxton et al., 2004).

When active neurons consume energy and oxygen, the surrounding capillaries dilate and the concentration of HbO and HbR in the tissue increases. As a result, the oxygen supply to the neuron also increases (Buxton et al., 1998; Friston et al., 2000). We can see this in more detail in Figure 2.3C.

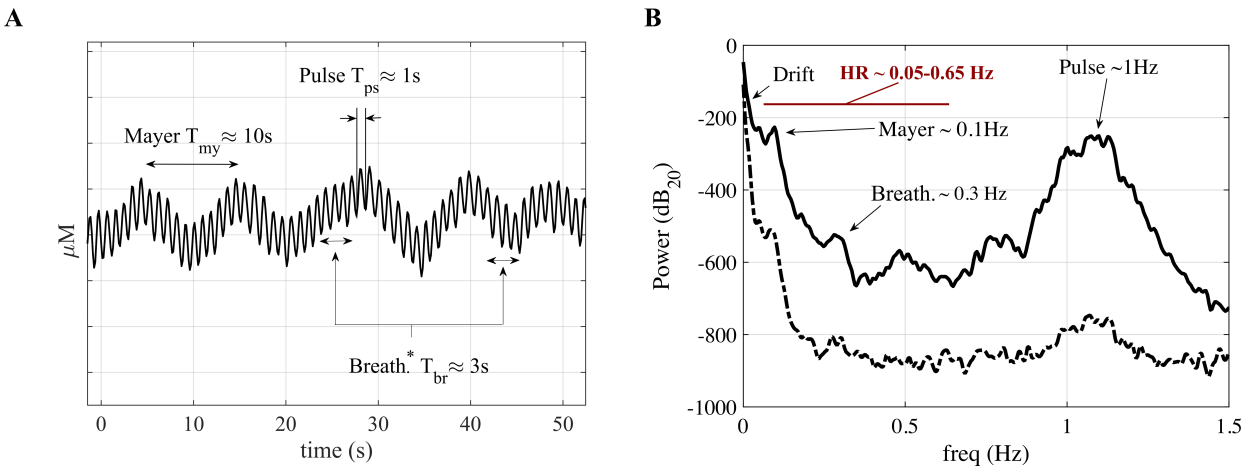


Figure 2.5: **A** Confounding sources of Hb changes in the time-domain. HbR not shown. Breathing* waves can hardly be identified as they have lower amplitudes than the rest of the artefacts. Their effect can be seen in the deformation of Mayer waves. **B** Ibidem in the frequency domain for HbO (continuous) and HbR (dashed).

First, oxygen's metabolic consumption triggers a physiological response in the surrounding vessels, which start dilating: the vasodilatory response. Consequently, the blood flow increases and, with it, the concentration of HbO and HbR supplied to the tissue. This physiological process's end goal is to supply the metabolically active neurons with a constant exceeding supply of oxygen. The changes in Hb concentrations after neuronal activity are known as the HR and are used as a biomarker of cortical brain activation. Figure 2.4A shows the different parameters that define the morphology of the HR.

In some cases, the HbO response starts with a slight and fast decrease, corresponding to the early oxygen depletion, just before overshooting (Fig. 2.4A and B). Then, once the neural activity ceases, the HbO goes back to the baseline level with an undershoot (Fig. 2.4A). The HbR response is less studied, but it is accepted that it decreases following neuronal activation and is negatively correlated with HbO (Cui et al., 2010). Although the morphology of the response (Tzourio-Mazoyer et al., 2002) and its dynamical causes have been well studied (Buxton et al., 1998; Friston et al., 2000) a complete physiological model of the HR is still lacking (Vanzetta and Grinvald, 2008; Figley and Stroman, 2011).

Since most filtering techniques are based on the (partial) knowledge of the HR, this lack of physiological understanding can limit fNIRS applications (Zhang et al., 2009). We will go back to this in Chapter 4 and demonstrate how model-free and data-driven techniques like DL can present some advantages over physiological model-based approaches in detecting functional HRs.

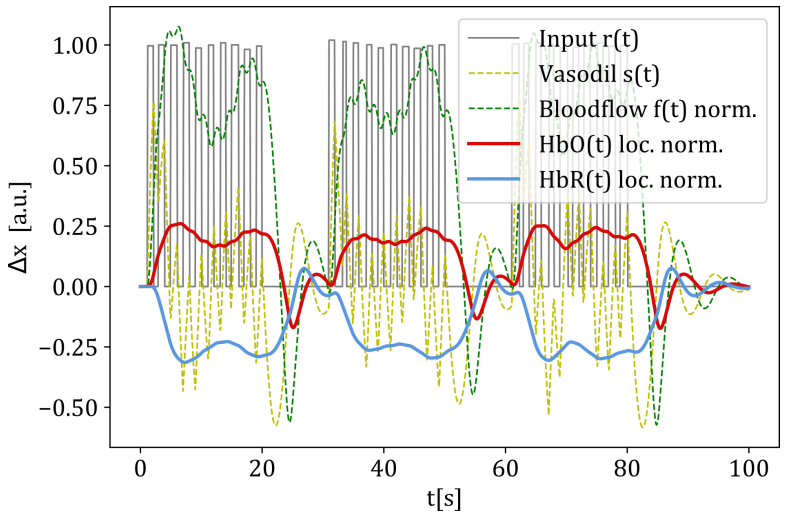


Figure 2.6: Simulated non-stationary haemodynamic responses. An neuronal metabolic demand (input, $r(t)$ in grey) causes a vasodilation ($s(t)$ in yellow) and consequently an increase in blood flow ($f(t)$, green) and HbO and HbR local concentrations (red and blue, respectively). The increase in HbO provides enough oxygen to the demanding neuron reaching an equilibrium. Measures are normalised. The HR model in Scarpa et al. (2010) was used to generate these HRs. In particular, the input driving the HR ($r(t)$, grey) was modelled to have a fast on/off switching dynamic with some amplitude variability. As a result, the HbO and HbR responses show some non-stationarities. In particular, we can identify two bumps in the first HbO response. In the second one, the first bump is mostly attenuated, while in the third, it reappears again. In some cases, fixed canonical HRs are used to filter out the noise. As we depict here, this could potentially lead to the loss of information since more realistic HRs show more variability. In other cases, models are used to create priors on expected HRs for given inputs, as shown here. In these cases, the input $r(t)$ is assimilated as the behaviour under study. However, the relationship between a specific behaviour and the actual $r(t)$ input driving the HR is unknown and probably more complex than the simple identity assumed in this kind of models, which is particularly relevant for cases where the HR is affected, e.g. after stroke and concussion.

fNIRS limitations

Changes in HbO and HbR are not only the result of neurovascular coupling but also heart beats (~ 1Hz), breathing (~ 0.3Hz) and vascular reflexes known as Mayer waves (~ 0.1Hz) (Fig. 2.3A) (Kirilina et al., 2012). The band-width of the HR is considered to be between $0.005 - 0.65$ Hz, overlapping with drift, breathing, and Mayer components (Fig. 2.5), which leads to the reduction of SNR of the HR and a loss of specificity of the spatio-temporal distribution of HRs in the cortex (Tachtsidis and Scholkmann, 2016; Yücel et al., 2016).

Several filtering techniques have been proposed to reduce these confounding sources' negative effect (Table 2.1). However, these methods are not flexible enough to cope with the nonlinearities, the non-Gaussianity and the partially unknown characteristics of the HR (Huppert, 2016). In particular, the Statistical Parametric Mapping (SPM) method (Ye et al., 2009), General Linear Model (GLM) (Abdelnour and Huppert, 2009) and wavelets (Jang et al., 2009) require the experimenter to assume an HR morphology before applying the filtering. This makes these approaches sensitive to biases in experimenter decisions. Other methods, like dynamic modelling (Friston et al., 2003), bayesian

filtering (Scarpa et al., 2010) and particle filters (Croce et al., 2017), compute the HR a priori, based on physiological models (Buxton et al., 1998, 2004). However, these physiological models are driven by the input neuronal activity, $r(t)$, which is the unknown we depart from and would like to reveal (Fig. 2.6). These methods assume that some mensurable aspect of the behaviour under study is equivalent to the neuronal input driving the HR and use it to model an expected HR.

In both cases, the expected HR is used to filter out the variability in the recorded signal that cannot be explained by the model (Friston et al., 2003; Ye et al., 2009). The problem is that these methods could be biased when studying diseases, leading to abnormal HRs; or behaviours that modulate the HRs in ways that are not captured by the expected responses (Bonakdarpour et al., 2007).

Other methods use short inter-optode distances (reducing the depth in the banana-shape) to record systemic variations of Hb at the scalp (Gagnon et al., 2011, 2012). The scalp haemodynamic activity is then regressed out of the brain haemodynamic activity. However, these approaches require additional instrumentation and are sensitive to the short distances used (Gagnon et al., 2012). As a result, an experimenter may end up rejecting brain activity as scalp activity.

We can turn these problems around with data-driven machine learning techniques, which allow the recorded brain signals to show us what is behaviourally relevant instead of injecting incomplete prior knowledge. In particular, DL techniques can automatically learn the relevant features in the brain signal and have been used to classify fNIRS signals (Hiwa et al., 2016; Trakoolwilaiwan et al., 2017). However, these approaches make use of black-box models that lack interpretability. As a consequence, they cannot be used to learn anything about the input brain signal.

In Chapter 4 we will show how machine learning, DL in particular, can be used to find relevant cortical signals associated with behaviour and to draw interesting neuroscientific interpretations from them.

N.C.	Method	Reference
688	SPM, GLM & gamma-functions	Ye et al. (2009)
208	Short optodes & state space	Gagnon et al. (2011)
207	GLM, gamma-functions & Kalman filter	Abdelnour and Huppert (2009)
171	Wavelets	Jang et al. (2009)
28	Bayesian Filtering	Scarpa et al. (2010)

Table 2.1: Top five relevant fNIRS processing methods by citation number (N.C.) (updated October 2020). All these approaches use a prior of the Haemodynamic Response (HR) to reduce noise. However, the HR is still not completely understood, and the use of such priors might lead to an underestimation of the response. SPM, statistical parametric mapping. GLM, general linear model.

2.2.2 Electroencephalography

Electroencephalography (EEG) was introduced by Berger in 1929 as a non-invasive way of measuring electrical brain activity (Haas, 2003). In contrast to fNIRS, EEG has a much higher time resolution (milliseconds) and provides direct measures of neuronal activity (Buzsáki et al., 2012).

Indeed, the mechanisms by which EEG records cortical neural oscillations are well understood and can be modelled accurately (Wang, 2010; Buzsáki and Wang, 2012). Neurons act like electrical dipoles. When hundreds of them fire synchronously in the cortex, their electrical fields constructively add up, generating a local field potential (Buzsáki et al., 2012). Local field potentials reach the outer scalp, where an electrical sensor can detect the voltage fluctuations.

EEG Limitations

EEG is characterised by a low spatial resolution (van den Broek et al., 1998; Leahy et al., 1998). Since the electrical fields travel in every direction[1] and the sensors capture the activity of thousands of neurons at the same time, it is hard to isolate the sources. Simultaneously, the scalp acts as a low-pass filter on the electrical fields, and higher frequencies are subject to lower SNRs (Ball et al., 2009; Buzsáki et al., 2012). Besides, head muscles, like the eye muscles, lips and chewing muscles, generate electrical potentials more powerful than EEG which leak into the EEG sensors. Other sources of external electrical activity (e.g. AC lights or electricity mains) can increase the ground level of noise and the measure's quality.

2.2.3 EEG and fNIRS

Compared to EEG, the muscular electrical activity does not interfere with fNIRS, which is a desirable characteristic when studying brain activity during motor behaviour experiments (Leff et al., 2011; Ferrari and Quaresima, 2012). In addition, it provides enough spatial resolution to distinguish between the sensory-motor area and motor and pre-motor cortices, all engaged in motor behaviour and very close to each other in the cortex (Leff et al., 2011). These two brain signals (fNIRS and EEG) have

[1] To be precise, volume conduction explains how electrical fields travel anisotropically through the biological tissues. Although it is possible to simulate how the electrical fields refract, it would require a precise definition of the tissue's anatomical configuration and refraction coefficients. There are simulations and models done at the population level but achieving this level of configuration details is not possible in practice for every subject.

very different physiological origins and offer complementary information about the brain. Hence, their combination seems promising in the field of non-invasive BCI.

In the next section, we will review the progress made in BCI force decoding. We will analyse the progress made in invasive BCI approaches and analyse the opportunities to improve non-invasive approaches. In particular, we will see how multimodal (fNIRS and EEG) BCIs have been used in the decoding of force.

2.3 Force decoding in BCIs: A systematic review

BCIs offer an alternative way to interact with our environment. They are especially relevant for people whose natural ability to interact physically has been lost or damaged. There has been a lot of progress in the decoding of kinematic variables in BCI (Carmena et al., 2003; Bansal et al., 2012; Sussillo et al., 2012; Sachs et al., 2015). In contrast, human force decoding is less explored. However, force is essential for safe and meaningful mechanical interactions (Branco et al., 2019b). In addition, force decoding can provide more generalisable BCIs decoders (Cherian et al., 2011) specially in changing dynamic conditions (Chhatbar and Francis, 2013). Even when the ultimate goal of a BCI is different from force decoding, understanding and exploiting signals that already exist in the brain might offer a more intuitive control than learning new ones (Flint et al., 2014).

In this review, we analyse invasive and non-invasive BCI methods for force decoding in humans and monkeys. The advance of force decoding in these directions could bring BCI applications closer to the rehabilitation, assistance and even replacement of force control. These advances acquire a greater significance in non-invasive BCIs since they provide an entirely safe, more affordable and accessible technology than its invasive counterpart. In this sense, the non-invasive modality could fulfil the needs of most patients for whom the benefits of an implantation surgery for invasive BCI do not overcome the risks or costs of the procedure.

2.3.1 Search criteria

The PubMed database was queried (last time on the 20th of October 2020) for publications containing the following terms in their title or abstract:

```
force AND decoding AND (‘‘brain computer interface’’ OR ‘‘brain machine interface’’)
                        -robot -auditory -mmg -learning -reach
```

This search resulted in 61 initial papers. Conference proceedings were not included unless they provided the first proof of concept of the use of a technique or method to decode force (we rejected 36 conference proceedings).

We also included relevant references consistently cited in the queried papers despite not appearing in the main search. In particular, these references included papers that did not specifically focus on force decoding but either provided interesting insights into the brain representation of force control or focused on decoding other motor control variables such as velocity or position (kinematics). Finally, we included $N = 38$ papers in this review, of which $N = 24$ specifically addressed the decoding of force or kinetics in the BCI context.

2.3.2 Overview of queried publications

A selection of 17 relevant force decoding publications is compared in Table 2.2. This table offers an overview of the aspects that characterise the relevant dimensions to take into account when analysing BCI methods for force decoding. Among these dimensions, we find that the queried publications span: (1) different subjects of study, from rats ($N = 3$) and monkeys ($N = 16$) to humans (in health and disease, $N = 21$); (2) different recording techniques: from invasive ($N = 28$) Microelectrode (ME) arrays, ECoG and DBS implants to non-invasive ($N = 21$) EEG and fNIRS; and (3) different decoding linear and non-linear techniques in the period between 2003 and 2020.

The applications vary from force decoding to some overlapping decoding of other motor control variables, such as movement or position. Other publications contained helpful information on the representation of force in the brain (for a review see Branco et al., 2019a), and we included them by their proximity to the force decoding topic.

Figure 2.7 depicts the evolution of the force decoding field in BCIs. Since one of the first force decoding demonstration in monkeys by Carmena et al. (2003), the field of continuous force decoding has been led by invasive experiments in these animals, moving 11 years later to invasive studies in humans and having only been non-invasively addressed by Paek et al. (2019).

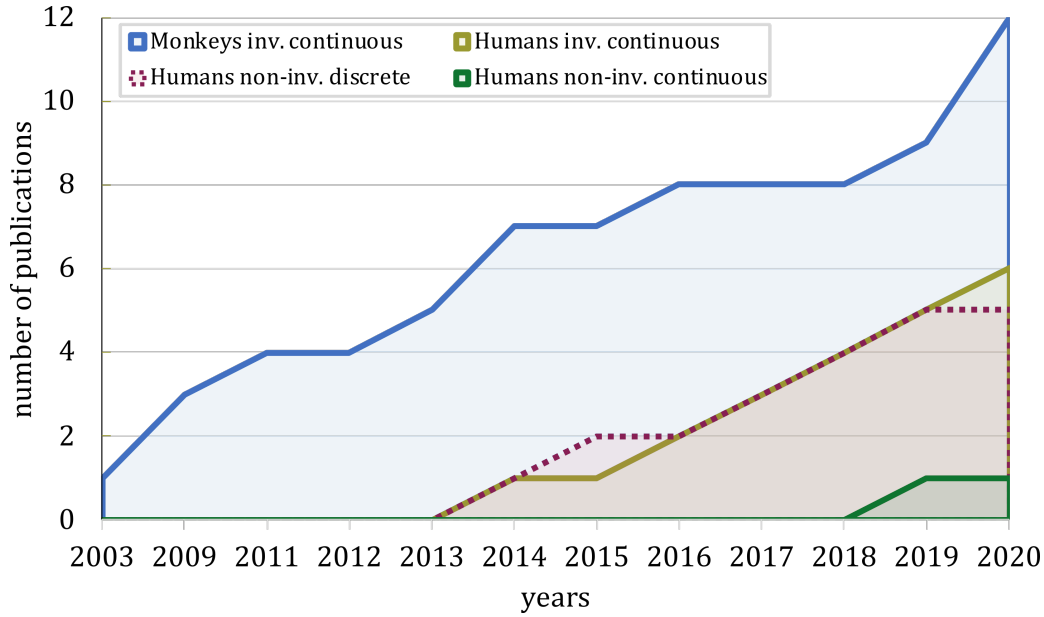


Figure 2.7: Accumulated force decoding publications for BCI from 2003 until October 2020. For the continuous decoding of force, i.e. the reconstruction of force trajectories from brain signals, much of the progress corresponds to invasive approaches for both humans (yellow) and monkeys (blue). Non-invasive approaches in humans mainly focus on decoding discrete force characteristics like the classification of force levels (red discontinuous). The non-invasive continuous decoding of force is, in comparison, a more recent field that requires further exploration (green).

	BCI control	Subjects	Rec. Tech.	Brain Area	Muscles	Feedback	Task	Method	Lag ms.	Depth ms.	Average Performance	Observations
Carmena et al. (2003)	Open l.	Monkeys	ME	M1, dPM, SMA, PP, S1	Arm, hand	Visual, robot	Reaching, Hand-grip, Reach/grasp	Wiener	0	1000	FVAF = 95% (hand grip force)	Force better decoded than kinematics.
Gupta and Ashe (2009)	Offline	Monkeys	ME	M1, dPM	Arm	Visual, propio.	Force fields	Wiener	20	420	r = 0.78 (arm EMG)	Kalman filter less appropriate for decoding.
Fagg et al. (2009)	Closed l.	Monkeys	ME	M1 (contra.)	Arm	Visual	Random pursuit	Wiener	0	1000	FVAF = 72% (arm torques)	Torques better for decoding.
Nambu et al. (2009)	Not causal	Humans	fNIRS (HbO)	M1	Hand fingers	Auditory, propio.	Isometric (finger pinch)	GLM	-4000	21000	FVAF = 46% (finger pinch force)	Contralateral activation (around C3/C4).
Cherian et al. (2011)	Open l.	Monkeys	ME	M1	Arm	Visual, propio.	Rand. pursuit (dyn. loads)	Wiener	0	500	R2 = 0.48 (arm EMG)	EMG better generalisation.
Chhatbar and Francis (2013)	Closed l.	Monkeys	ME	M1	Arm	Visual	Rand. pursuit (dyn. loads)	Wiener	0	100	r = 0.73 (arm torques)	Torque better BCI control.
Chen et al. (2014)	Offline	Monkeys	ECoG	SM (left)	Arm	Visual, propio.	Reach/grasp	Sparse GLM	0	na	r = 0.74 (reach/grasp)	High gamma band more informative.
Hao et al. (2014)	Offline	Monkeys	ME	dPM	Hand	Natural	Reach/grasp	SVM	0	na	Class. Acc. = 92.5% (4 grip classes)	Grasp-specific neurons. Hand-grip higher firing rates than finger pinch.
Flint et al. (2014)	Offline	Humans	ECoG	SM	Hand fingers	Natural	Isometric (finger pinch)	Wiener (3rd order)	0	20	FVAF = 60% (finger pinch force)	High gamma and time domain more info. Finger specific.
Hansen et al. (2014)	Offline	Humans	fNIRS (HbR), EEG	F, FC, T, C	Right foot	Visual, propio.	Isometric	LDA	0	15000	Class. Acc. = 76% (Bin. mov. detection) Class. Acc. = 81% (Bin. force 20,60%MVC)	EEG better in movement detection and fNIRS in force level decoding.
Ethier et al. (2016)	Closed l.	Monkeys	ME	M1 (contra.)	Wrist	Visual	Isometric	Wiener casc. (sigmoid)	0	500	MSE = 20 (wrist EMG)	Force more stable decoding.
Fu et al. (2016)	Offline	Humans	fNIRS, EEG	SM	None	Self-paced	Speed, force (imageries)	SVM	na	na	Class. Acc. = 74% (6 imageries)	Multimodal better classification.
Shah et al. (2017)	Offline	Humans (Parkinson)	LFP	STN	Hand	Self-paced	Isometric (hand grip)	Regularised Wiener casc. (3rd order)	0	1256	r = 0.67 (hand grip force)	High gamma more informative in STN.
Shah et al. (2018)	Offline	Humans (Parkinson)	LFP	STN	Hand	None	Isometric (hand grip)	Regularised Wiener casc. (3rd order)	0	256	r = 0.55 (hand grip force)	Non-linearities improve decoding.
Naufel et al. (2019)	Offline	Monkeys	ME	M1	Wrist	Visual, propio.	Isometric (dyn. loads), movement	Wiener (non-lin.), LSTM	0	500	FVAF = 70% (wrist force)	Non-lin. more robust decoding in changing dyn. loads.
Paek et al. (2019)	Offline	Humans	EEG	Cortex	Hand fingers, arm	Visual, propio.	Isometric (finger pinch), grasp/lift	GLM	0	250	r = 0.42 (finger pinch force) (5,10,15%MVC)	Potential non-lin. force/ERD relationship.
Ganzer et al. (2020)	Closed l.	Human (SCI)	ME	M1	Hand	Visual, augmented propio.	Reach/grasp	SVM, GLM	0	na	Class. Acc. = 86% (touch decoding)	Restore motor function and sense of touch.

Table 2.2: Comparison of queried studies specifically targeting force decoding. The table includes: the recording technology used (Rec. Tech.), the subjects under study, the recorded brain area, the muscles used during the task, the type of feedback received to accomplish the task, the type of task, the decoding method, the time horizon in milliseconds of the prediction ("lag", positive values indicate how far in the future the method predicts the force), the amount of history used in the decoding (depth), the average performance and key observations.

Before, non-invasive techniques focused on decoding discrete features of force, such as the differentiation between discrete levels of force (Wriessnegger et al., 2017) or between executions and imageries (Yin et al., 2015; Wriessnegger et al., 2017). Then, very recently, Paek et al. (2019) achieved moderate decoding performance of finger-pinch forces using EEG. Thus, it is clear that non-invasive approaches, particularly for the continuous decoding of force, are still emerging and require further exploration.

On a final note, some publications used EMG as a way to measure force. It is worth noting that decoding EMG entails a different challenge than decoding force. EMG signals can be less invariant to fatigue effects or muscle length. However, these parameters influence the force that the same muscle produces under the same EMG activity. Hence, the decoding of end-point forces can be considered more complex than that of EMG, which is closer to cortical brain activity (Cheney and Fetz, 1980; Fromm, 1983).

To shed light on the current situation of the BCI force decoding field, we formulate the following questions: (1) What techniques are used to record brain signals with force decoding purposes, and which brain areas are targeted? (2) What force-related information do brain signals carry? (3) How does force decoding depend on the force task? (4) Is the decoding of force generalisable across task or subjects? Finally, (5) what methods have been used to decode force?

We review and discuss the relevant information that the queried papers contain regarding these questions and offer an overview of the advances that have been made in the last two decades. The analysis of these questions allows us to identify gaps that have been less or not explored. These gaps bring new opportunities in advancing force decoding in non-invasive BCI with applications in the rehabilitation, assistance and even replacement of force control.

2.3.3 What techniques are used to record signals for force decoding and what are the brain areas targeted?

To decode force from brain signals, we need to know where force-related information is available in the brain and how can it be accessed. Most brain activity recording techniques have been used for force decoding purposes. Invasive techniques, from more to less invasive, include: DBS implants (Fischer et al., 2017; Shah et al., 2017, 2018), ME arrays (Carmena et al., 2003; Gupta and Ashe, 2009; Fagg et al., 2009; Cherian et al., 2011; Bansal et al., 2012; Chhatbar and Francis, 2013; Aggarwal et al., 2013; Hao et al., 2014; Sachs et al., 2015; Ethier et al., 2016; Ganzer et al., 2020) and ECoG (Pistohl et al.,

2013; Chen et al., 2014; Flint et al., 2014, 2016; Ryun et al., 2017; Branco et al., 2019a). Non-invasive techniques include: fMRI (Ganesh et al., 2008), EEG (Paek et al., 2019), fNIRS (Nambu et al., 2009; Wriessnegger et al., 2017) or both, EEG and fNIRS, combined (Hansen et al., 2014; Yin et al., 2015; Fu et al., 2016).

Regarding non-invasive techniques, one desired aspect in BCI applications is the wearability, i.e. how easily the recording system can be transported and worn, with enough power autonomy to move around freely (Casson et al., 2010). EEG and fNIRS are portable, in some cases wearable, and have electronics that allow for miniaturisation (Casson et al., 2010; Pinti et al., 2018, e.g.) while fMRI requires scanners generating powerful magnetic fields that cannot be worn by subjects. Besides, subjects are confined to a reduced space while inside the fMRI scanner, which constrains their movement range. Therefore, the variety of tasks they can perform inside the scanner is limited by insufficient space.

At the expense of breaking biological tissues (skin, skull and pia layers), invasive systems can target specific brain areas and record information very close to its source. ME arrays provide the highest spatial resolution at the cortical level even reaching the single neuron resolution (Carmena et al., 2003; Gupta and Ashe, 2009; Bansal et al., 2012). However, the single-neuron resolution does not seem to bring any advantage compared to neuronal ensemble activities Carmena et al. (2003), eliminating the need to sort the neuronal ensemble spikes. Nonetheless, the amount of neurons in the ensemble is important, with more neurons providing better decoding Gupta and Ashe (2009).

Local Field Potentials (LFPs) (the near-field sum of the electrical neuronal ensemble activities) can also be measured with ME arrays going a step further down in spatial resolution. Although LFPs are less informative than spikes in force trajectory reconstruction (Bansal et al., 2012; Aggarwal et al., 2013), they are better at detecting and predicting other force characteristics like the onset and end of force generation (Aggarwal et al., 2013). In LFPs, force related information is distributed in frequency bands. In particular, LFP amplitudes in low frequencies (15 – 22 Hz) and LFP power in high frequencies (100 – 175 Hz) provide better force trajectory reconstruction (Aggarwal et al., 2013).

ECoG, the invasive technique with lowest spatial resolution, also shows a similar spectral distribution of force related information (Pistohl et al., 2013; Chen et al., 2014; Flint et al., 2014; Branco et al., 2019a). Similarly, low-frequency amplitudes, also known as Local Motor Potentials (LMPs), and high-frequency power show better decoding followed by the beta band (Pistohl et al., 2013).

	bandwidth [Hz]	Measure	Rec Tech.	Force decoding
Cortical Spikes		Firing rates	ME	Hand-grip force (Carmena et al., 2003), arm torques, reach, grasp. Grasp and sense of touch restoring in SCI person (Ganzer et al., 2020).
STN-LFP		Amp. [V]	DBS	Isometric hand grip force,
LMP	< 3	Amp. [V]	ECoG	Finger pinch force (Flint et al., 2014), hand kinetics and kinematics (Chen et al., 2014; Flint et al., 2016; Ryun et al., 2017)
			EEG	Higher voltage changes 100 ms before power-grip in posterior contralateral channels (Iturrate et al., 2018).
Delta δ	1 – 4	Pow. [dB]	ECoG	Power decreases during motor control. Better prediction of continuous force (Flint et al., 2014, 2016).
			EEG	Power decreases during motor control.
Theta θ	4 – 8	Pow. [dB]	ECoG	
			EEG	Moderate linear reconstruction of forces, strongest posterior correlation in isometric (Paek et al., 2019).
Mu μ (alpha α)	8 – 13	Pow. [dB]	ECoG	Movement detection (ERD), low muscle-like modulation. Worse decoding of continuous force better movement/rest differentiation (Flint et al., 2014, 2016).
			EEG	Power decreases during motor control (ERD). More informative (Iturrate et al., 2018; Hayashi et al., 2019) Moderate linear reconstruction of forces and stronger force correlations in posterior channels (Paek et al., 2019). Higher ipsilateral power suppression during precision grip (Iturrate et al., 2018).
Beta β	13 – 30	Pow. [dB]	ECoG	Strong power suppression (~ 20 Hz) (Flint et al., 2014)
			EEG	Power decreases during motor control (ERD). More informative (Iturrate et al., 2018; Hayashi et al., 2019). Strongest posterior correlation in isometric, moderate linear reconstruction of forces (Paek et al., 2019). Higher ipsilateral power suppression during precision grip (Iturrate et al., 2018).
Low gamma γ_L	30 – 50	Pow. [dB]	ECoG	Power decreases during motor control (ERD).
			EEG	Less informative (Iturrate et al., 2018; Hayashi et al., 2019). Low linear reconstruction of forces, frontal correlations in isometric (Paek et al., 2019).
High gamma γ_H	> 70	Pow. [dB]	ECoG	Grasping, flexor/extensor onsets (Flint et al., 2014), somatosensory feedback (Ryun et al., 2017; Branco et al., 2019b). Higher power for higher force production (Flint et al., 2014; Chen et al., 2014). More predictive of movement? (Flint et al., 2016).
			EEG	Less informative (Iturrate et al., 2018; Hayashi et al., 2019). Low linear reconstruction of forces, frontal correlations in isometric (Paek et al., 2019).
HbO/HbR	0.01 – 0.3	Conc. [mol]	fNIRS	HR modulated by speed (Obrig et al., 1997; Kuboyama et al., 2005) and intensity of force (Shibuya et al., 2008; Derosière et al., 2014; Wriessnegger et al., 2017). Discrete levels of foot force (20 and 50%MVC) (Hansen et al., 2014) and hand force imageries (Fu et al., 2016). Acausal reconstruction of finger pinch force (Nambu et al., 2009).

Table 2.3: Force related information per brain signal including spikes, Local Field Potentials (LFPs), ECoG and EEG frequency bands and fNIRS in the queried literature. Although the force has been decoded from invasive techniques, EEG and fNIRS have been less explored in the decoding of forces.

Regarding the brain areas that should be invasively targeted for force decoding, the publications show varying preferences. Interestingly, invasive studies show that force can be indistinctly decoded from different areas involved in motor control, such as M1, S1, Premotor ventral cortex (PMv) and Premotor dorsal cortex (PMd) (Carmena et al., 2003; Gupta and Ashe, 2009; Bansal et al., 2012). In particular, Carmena et al. (2003) showed that the monkeys' M1 and Posterior parietal cortex (PP) were better at decoding hand-grip force. The force predicted from these areas accounted for a maximum of 85% of the original force variance.

In contrast, the SMA, PMd and S1 were better at decoding kinematics. Carmena et al. (2003) also pointed that the contralateral M1 and S1 could be closer to motor control. Instead, the PMv, SMA and ipsilateral M1 could be more involved in learning visuomotor spatial relationships. The topographical distribution of force related neural control also varies with time, and different neurons show preferential activation at specific times before the force onset, some firing at the onset (Gupta and Ashe, 2009).

Besides cortex, it has also been possible to decode force from deeper brain structures like the Subthalamic Nucleus (STN) (Shah et al., 2017, 2018). The variety of brain areas with force decoding capabilities is not surprising since different brain areas simultaneously use the same force-related information with different aims (Scott, 2016).

Non-invasive technologies cannot reach the single neuron resolution and have, overall, a coarser spatial resolution and lower SNR (Ball et al., 2009; Buzsáki et al., 2012). Nonetheless, fNIRS has higher spatial resolution than EEG and has been used to specifically target the sensorimotor cortex (Nambu et al., 2009; Hansen et al., 2014; Yin et al., 2015; Fu et al., 2016) while EEG, in practice, is used to target the whole cortex (Iturrate et al., 2018; Paek et al., 2019; Hayashi et al., 2019). In terms of the characteristics of the signals used, EEG studies primarily focus on the spectral power and Local Motor Potential (LMP) (Iturrate et al., 2018; Paek et al., 2019; Hayashi et al., 2019). Alternatively, fNIRS studies exploit different temporal and amplitude features in HbO (Nambu et al., 2009), HbR (Hansen et al., 2014) or the difference between them (Yin et al., 2015; Fu et al., 2016). However, other combinations of HbO and HbR remain unexplored probably due to the complexity and lack of understanding of the HR during force control (Derosière et al., 2014; Shibuya et al., 2016; Wriessnegger et al., 2017).

Altogether, we can observe a widespread use of invasive techniques despite their complexity and the risks they pose to health but mostly carried out in animals instead of humans (Fig. 2.7). As a positive

note, invasive research provides a clear and invaluable understanding of where and what form of force-related activity can be detected in the brain. In comparison, non-invasive technologies tend to require extensive exploration of signal features to provide force-related information. This situation is clearer in fNIRS, a newer technique for the decoding of force (Wriessnegger et al., 2017). In both cases, non-invasive techniques depend on heavy feature engineering. In this context, the use of methods that automatically learn these force related features might accelerate the development of non-invasive BCI for force decoding.

2.3.4 What force related information do brain signals carry?

Different recording techniques can capture different aspects of the same underlying brain activities. To better understand what aspects of force can be decoded from each brain signal, literature often points at the division of force generation characteristics in two levels: low-level or muscle-like characteristics and high-level characteristics (Cherian et al., 2011; Chhatbar and Francis, 2013; Aggarwal et al., 2013; Chen et al., 2014; Hao et al., 2014; Flint et al., 2014, 2016).

The low-level characteristics of force are related to the force precise control and are also called muscle-like characteristics due to their proximity to muscle activity control (Cherian et al., 2011). Carmena et al. (2003) showed that these muscle-like characteristics, i.e. fine-grained reconstruction of force trajectories, could be decoded in monkeys better than kinematic variables (velocity and position) in specific brain areas (M1 and PP).

In contrast, high-level force characteristics refer to more abstract ones, such as the preparation, the anticipation, and the posture before force generation. It is worth noting that low-level force characteristics include the detailed decoding of EMG activity and the decoding of end-point forces, but each entails different levels of complexity. In particular, muscle activity is better decoded than end-point forces (Chhatbar and Francis, 2013). The relationship between muscle activity and end-point force depends on how fatigued the muscle is and the length and the configuration of the muscle (Cheney and Fetz, 1980; Fromm, 1983). These are unknown correlated variables that should also be decoded from the brain signals, thus increasing the problem's complexity. As a result, end-point forces can be decoded better if the anatomical configuration and the arm dynamics are known and incorporated in the decoding scheme (Chhatbar and Francis, 2013).

In general, but not exclusively, muscle-like force characteristics have been decoded from spiking or neu-

ronal ensemble activities in humans (Bouton et al., 2016; Ganzer et al., 2020) and monkeys (Carmena et al., 2003; Gupta and Ashe, 2009; Fagg et al., 2009; Cherian et al., 2011; Chhatbar and Francis, 2013). Two Degrees of Freedom (DoF) of force have been disentangled from spiking activities in monkeys (Chhatbar and Francis, 2013; Ethier et al., 2016), even though these firing rates do not seem to be highly modulated by different tasks (Naufel et al., 2019). The ability to simultaneously decode different DoF from the same brain signals is an interesting BCI challenge that would allow a much richer control of force. In particular, the successful decoding of different DoF would enable the application of force in different directions or the simultaneous use of different prosthetic limbs, further expanding the application of BCI in daily interactions.

While it is clear that detailed muscle-like characteristics can be decoded from single neurons or neuronal ensembles –provided that they capture enough units (Carmena et al., 2003)– if muscle-like characteristics can be decoded from invasive techniques with lower spatial resolution, like ECoG, is still a subject of debate. ECoG is an invasive technique that provides a similar temporal resolution than ME arrays but a smaller spatial resolution. In ECoG, electrodes placed millimetres away from the cortex but still inside the skull measure the collective activity of thousands of neurons.

There is evidence showing that ECoG carries information to decode finger-pinch forces (Flint et al., 2014) and hand kinetics and kinematics (Ball et al., 2008; Chen et al., 2014; Flint et al., 2016; Ryun et al., 2017), including kinematics with 2-DoF (Pistohl et al., 2008). However, the same authors show that ECoG is better at decoding high-level force characteristics, like force or movement onsets, than it is at decoding muscle-like ones (Ball et al., 2008; Flint et al., 2014; Chen et al., 2014; Flint et al., 2016; Ryun et al., 2017).

These differences are made more apparent depending on the frequency band used for the decoding (Table 2.3). In particular, high-gamma bands (> 70 Hz) reflect spiking activity in monkeys (Ray et al., 2008) and rats (Yazdan-Shahmorad et al., 2013) and are often used in ECoG for muscle-like decoding purposes (Flint et al., 2014; Chen et al., 2014; Flint et al., 2016). In the high-gamma band ECoG power increases at the onset of force application (Flint et al., 2014) for flexor and extensor grasp muscles (Flint et al., 2016). Furthermore, high-gamma power is higher for higher force production in the force onset (Flint et al., 2014; Chen et al., 2014).

Nonetheless, there is evidence that high-gamma power primarily represents somatosensory feedback rather than efferent motor commands. This sensory feedback reaches S1 tenths of milliseconds before

the same information reaches M1 on a lower scale (Ryun et al., 2017; Branco et al., 2019b). Remarkably, this kind of somatosensory feedback has proven essential to re-establish motor control in an SCI patient for whom three levels of grip force could be decoded using these signals (Ganzer et al., 2020). More recent evidence shows that when the contractions are isometric, i.e. not accompanied by movement, high-gamma power is not better than lower frequencies at decoding muscle-like characteristics of force, and high-gamma does not even show a sustained[2] response (Flint et al., 2016; Branco et al., 2019b). Therefore, high-gamma might be more valuable in decoding force when accompanied by movement (Flint et al., 2014).

The information in other bands with lower frequencies (< 100 Hz) is still also a subject of debate. In the low ECoG frequencies, amplitudes (< 3 Hz) (local motor potential, LMP Schalk et al., 2007) have shown better force decoding performance than the high-gamma power (Flint et al., 2014). In contrast, power in lower frequencies ($4 - 35$ Hz, mu, beta and low gamma) is higher during rest and suppressed during active grip, especially in beta (~ 20 Hz) (Flint et al., 2014). Hence, mu-beta power decrease during motor execution (Event Related Desynchronisation, ERD Pfurtscheller and Lopes Da Silva, 1999) is not exclusive of kinematics but is also present during isometric force production (Flint et al., 2014, 2016). Nonetheless, mu-beta ($7 - 20$ Hz) Event Related Desynchronisation (ERD) provides worse decoding of continuous force than the delta band and better differentiation of movement versus rest (Flint et al., 2014, 2016), and is also more modulated by the kinematics than by the kinetics of the movement (Flint et al., 2016).

Some of these spectral cortical LFP activities are also detected in deeper brain structures like the STN. Namely, Fischer et al. (2017) and Shah et al. (2017) showed that gamma power activity in the STN is modulated by intended grip force even before the force onset and even when the force was only imagined. Hence, they showed that LFP activity is not only related to sensorimotor feedback at the STN level. Interestingly, lateralisation of activity in the STN was only observed during actual force generation but not imagined force (Fischer et al., 2017; Shah et al., 2017). As in ECoG, the high gamma band ($300 - 500$ Hz) had the highest decoding power in muscle-like decoding, and beta ($13 - 30$ Hz) and low gamma ($55 - 90$ Hz) could better predict force onsets 1.3 s in advance, but all the features together had the best force reconstruction power (Shah et al., 2017).

[2] In single-neuron studies, neuronal responses are categorised in tonic and phasic. Tonic responses are sustained during movement or force production and reflect the amplitude or progression during a motor control task. Phasic responses happen in the transition of motor control states and signal a change of specific control modes, for instance, when starting or stopping a movement.

Non-invasive techniques have also shown potential for decoding force but have been less explored compared to invasive techniques, in part, due to their lower SNR (Ball et al., 2009; Buzsáki et al., 2012). On the one hand, EEG has been used to decode EMG activity (Hayashi et al., 2019) but has shown limited reconstruction ($r \approx 0.42$) of forces (Paek et al., 2019). Indeed, as in the invasive case, targeting EMG avoids the decoding of the non-linear transformations that turn the muscular activity into end-point forces simplifying the decoding (Úbeda et al., 2018). In contrast to ECoG studies, in EEG the most informative bands for force decoding are the delta band ($1 - 4\,\text{Hz}$) and the alpha ($8 - 13\,\text{Hz}$) and the beta ($13 - 30\,\text{Hz}$) bands (Iturrate et al., 2018; Paek et al., 2019; Hayashi et al., 2019). In addition, although higher frequency bands (i.e. gamma bands) provide less force decoding information (Paek et al., 2019), they are still detected with EEG during movement (Ball et al., 2008).

On the other hand, Nambu et al. (2009) could also decode a smoothed version of force using a non-causal decoder (including 3 s of future information) using fNIRS. However, the relationship of the Haemodynamic Response (HR) with force seems more complex than previously described (Dettmers et al., 1995; Wexler et al., 1997; Thickbroom et al., 1998) and seems to be modulated by the speed (Obrig et al., 1997; Kuboyama et al., 2005) and the intensity and duration of force generation (Shibuya et al., 2008; Derosière et al., 2014; Shibuya et al., 2016; Wriessnegger et al., 2017). In particular, it has been shown that higher forces (40% of the Maximum Voluntary Contraction, MVC) lead to faster Hb peaks (~ 3 s before) than lower forces (20% Maximum Voluntary Contraction (MVC)) (Wriessnegger et al., 2017). However, the temporal properties of fNIRS are rarely explored and changes in HR amplitudes at predefined times are preferentially used in the analyses (e.g. Shibuya et al., 2008; Derosière et al., 2014; Hansen et al., 2014; Shibuya et al., 2016).

EEG and fNIRS have been used together (also known as multimodal or hybrid systems) to decode force characteristics. It has been shown that their combined use improves the classification of discrete force characteristics (Hansen et al., 2014; Yin et al., 2015). In particular, fNIRS seems to provide higher classification of force levels: 80% classification accuracy distinguishing 20%MVC from 60%MVC for fNIRS compared to 61% for EEG. In contrast, EEG is better at distinguishing movement from rest, 74% accuracy for EEG vs. 65% for fNIRS (Hansen et al., 2014). However, multimodal systems for force decoding are not very prevalent in the literature and have not been yet explored for the continuous decoding of force.

Finally, in both the invasive and the non-invasive case, a more recent and exciting way of understanding the division between high-level and muscle-like characteristics of force in the brain is beginning to

emerge. As we have seen, high-level force characteristics can be represented by brain signals at the neuronal population levels like during force preparation in the ECoG gamma band (e.g. Flint et al., 2014; Chen et al., 2014) or the grasp and release in the mu-beta band (e.g. Pfurtscheller and Lopes Da Silva, 1999; Flint et al., 2014). The evidence suggests that these neuronal population responses signal abrupt transitions in the activity at the single neuron level (Aggarwal et al., 2013; Hayashi et al., 2019) that vary, at least, with the force level and that require non-linearities to account for this transition (Aggarwal et al., 2013; Naufel et al., 2019). These abrupt changes in brain activities point at discontinuities or non-linearities in the brain function as it switches from one control mode (e.g. reaching) to another (e.g. grasping). If the brain signals modulate the force (or vice versa) in non-linear ways, then non-linear methods would be more appropriate to decode force from brain signals, as Naufel et al. (2019) recently showed in the invasive case. However, these non-linear methods have yet to be explored for the continuous decoding of force from non-invasive signals.

2.3.5 How does force decoding depend on the motor task?

Force-movement covariation During motor control, the force generated by the muscles results in movement through lever-like mechanisms. In the presence of an obstacle impeding the end-point movement, the muscles maintain their length constant, and the contraction causes an increase of pressure at the end-point (isometric contraction). Conversely, when movement is allowed, the task is characterised by the generated force and the velocities and positions of the moving body part under study. These forces and velocities can co-variate due to the mechanical configuration of the body and the task, difficulting the separation of force decoding (i.e. the ability to decode force itself) from the decoding of the covarying velocities and positions, which represents an easier task (Fagg et al., 2009).

In the invasive case, Carmena et al. (2003) first showed that the force could be decoded from a monkey's brain signals, whether or not it was accompanied by movement. The main difference between motion and force decoding was found across different brain areas: the M1 and the PP provided the best accuracy in decoding forces while the SMA, PMd and S1 provided better performance in decoding kinematics.

The role of somatosensory feedback In particular, S1 is essential in decoding sensory feedback (Ryun et al., 2017; Branco et al., 2019b) which suggests that force decoding could be dependent on the level of feedback information available to perform the decoding. The way a task modulates the

feedback could explain, to some extent, why some tasks are better or worse decoded than others. First of all, the feedback during movement is likely richer than during isometric contractions. This would explain the lower performance of S1 decoding grip forces in Carmena et al. (2003) and better performance of S1 compared to M1 in decoding the movement types (Ryun et al., 2017). Furthermore, the amplification of this kind of feedback, from sub-perceptual to perceptual, from S1 brain signals enabled a SCI patient to restore his grasp (Ganzer et al., 2020). The latter introduces the idea that a simple on-off controller model is enough to restore grasp, at least to a pragmatical extent. The on-off controller in Ganzer et al. (2020) shows that M1 can be voluntarily modulated depending on information present in S1. Although this work does not show that M1 directly encode force in humans with quadriplegia, it suggests that force control could be closer to M1 activity modulated as a function of the remaining desired effort (Cheney and Fetz, 1980; Fromm, 1983) which is potentially represented in S1.

Indeed, in the STN, a processing centre closer to the spinal cord and the muscles, gamma-beta activity is significantly correlated with three different force levels of force before they were executed and, even, when their execution was prevented (Fischer et al., 2017). This further suggests that the control of contractions might be relative to the current level of contraction. From a resting state, the remaining desired effort is maximum, and it decays as force increases. Ultimately, the intimate dependence in hand force decoding with somatosensory feedback indicates that the decoding of force relies on the efferent command to the muscle, which at the same time depends on the afferent signals from the muscle or skin to the brain, in a closed-loop fashion (Scott, 2016).

Additional evidence suggests that decoding muscular activity (with our without movements) is more straightforward than decoding end-point forces (Cherian et al., 2011; Chhatbar and Francis, 2013; Ethier et al., 2016). This evidence might also shed some light on how force generation and its decoding depend on the somatosensory feedback. Namely, decoding muscular activity does not consider all the non-linear transformations that the contraction of muscles undertake until being transformed in end-point forces (Úbeda et al., 2018). Hence, the decoding of muscular activity might require less information on the end-point's current situation, represented in the brain through sensory feedback to S1. As Carmena et al. (2003) showed, S1 was better at decoding kinematics, which might represent how the brain encodes the consequences of the forces it commands to apply. For example, Fagg et al. (2009) shows that torques from the elbow are more complicated to decode than those of the shoulder, suggesting that the more complex the mechanical conversion (the elbow torques are more variable and

depend on the shoulder ones), the worse the decoding. Nonetheless, even if force decoding is more dependent on sensory feedback, it has been shown that this sensory feedback can be retrieved even in cases where a complete SCI blocks the conscious perception (Ganzer et al., 2020).

In the non-invasive case, there is also some evidence showing that force is better decoded when accompanied by movements. In particular, the delta EEG band has been found to have more significant correlations with decoded forces when movements accompanied them than in the isometric case (Paek et al., 2019). Moreover, the movement and isometric task also led to different areas activated for different frequency bands. During force generation with movement, frontal and dorsal cortical areas were bilaterally more correlated with delta than during the isometric task. For the isometric task, theta, alpha, and beta had bilateral dorsal activations, which were more frontal for gamma. In contrast, the power in the beta band was more relevant for force onset detection as in ECoG (Flint et al., 2014, 2016).

Muscle selectivity The presence or not of movement is not the only factor that might intervene in force decoding. Since each muscle has a specific representation in the brain (Penfield and Boldrey, 1937; Flint et al., 2014; Úbeda et al., 2018), force decoding can also be selective of the muscles engaged in a task and their number. For example, the decoding of hand forces, having the most extensive cortical projections, might require a higher number of neurons controlling all the subtle forces that the hand can produce. Simultaneously, this higher representation might make the hand-related information more available, increasing the signal to noise ratio. Even in the case of symmetrical limbs (e.g. the arms), differences in decoding might arise as a consequence of the dominance of one limb over the other. For example, in right handed people, the right hand force seems to be better reconstructed than the left hand force (FVAF right/left = 0.5/0.42) (Nambu et al., 2009). Nonetheless, except (Nambu et al., 2009) the queried studies only target force generated with the dominant limb.

Furthermore, muscle selectivity in force decoding can appear even for muscles controlling the same limb. For example, within the same hand, the release of a grasp seems to elicit a 70% less ECoG power than its contraction, especially in the isometric case (Branco et al., 2019a). Some of these differences might be, to some extent, explained by differences in the cortical topography of the muscles involved. In particular, hand extension and contraction selectively engage sparse and intermingled cortical areas (Ganesh et al., 2008). These cortical areas influence the decoding of hand forces to the finger level (index or little finger to thumb pinches) with different fingers leading to shifts in the

most relevant positions for the decoding (Flint et al., 2014). Overall, cortical muscle representation is one of the factors that influence the decoding of force. We can expect that the more intermingled the representation of muscles is at the cortical level, the worse the decoding will be for the sensing techniques with coarser spatial resolution.

Another factor influencing the decoding of forces is the timing with which brain responses orchestrate force generation (Chen et al., 2014). For example, Hao et al. (2014) found that some neurons in the PMd fire selectively to the kind of grasp executed with specific locations and timings. In particular, Hao et al. (2014) show that firing rates in the PMd of monkeys are predictive of the type of grasp 200 to 300 ms before the grasp is secured. After securing the grasp, the neurons' specificity decreased, and higher firing rates corresponded to power grasps during the hold period.

Muscle coordination Other studies put the covariation of kinetics and kinematics in practice and decode only the two or three Principal Components (PCs) explaining most of the variability of the kinetic and kinematics variables defining a movement (Chen et al., 2014; Flint et al., 2016). In particular, PCs reflecting fine-grained characteristics of movement are better decoded by the ECoG gamma band than lower frequency bands (Flint et al., 2016). Nonetheless, as we mentioned earlier, the nature of this analysis makes it challenging to establish whether the ECoG gamma band is good at decoding force by itself; or this result is just a consequence of forces and kinematics being decoded simultaneously.

As we have seen, the decoding of force is sensitive to the muscles engaged in a task, whether or not they are allowed to generate movement, and how their activity is coordinated. Since all these factors define the type of task, we can expect that different tasks, highlighting different motor control abilities, show different decoding performances. For example, we can think about the abilities involved in picking a needle up from a flat surface or lifting a bottle of water to drink. The first case requires the precise control of two fingers, the thumb and the index, in a pinch like movement. Finger-pinch force is decoded in Nambu et al. (2009); Flint et al. (2014) and Paek et al. (2019). Instead, when we lift a bottle full of water, we use the whole hand to grasp the bottle. This kind of grasp or grip is often referred to as power-grip since it allows for greater forces to be applied. Power-grips are studied in (Carmena et al., 2003; Shah et al., 2017, 2018). Each of these grips requires a different set of muscles, forces and precision, and all these characteristics might be decoded differently from the brain.

Regarding the decoding of different levels of force, the STN seems to have less predictive power for lower than higher levels of force (Shah et al., 2017). This situation seems to be similar for EEG where

low forces tend to be overestimated (Paek et al., 2019). Nonetheless, evidence suggests that EEG is not linearly modulated with the level of force (Paek et al., 2019; Hayashi et al., 2019). For example, (Hayashi et al., 2019) shows that alpha and beta bands have multilevel force decoding capabilities and that the main difference lay in the cortical regions engaged for different levels. In their work, Hayashi et al. (2019) provide limited evidence that the alpha-beta band of bilateral frontoparietal regions encodes different levels of force while the contralateral parietotemporal regions seemed to differentiate better the force onsets (Flint et al., 2014, 2016). However, the results of Hayashi et al. (2019) should be taken with caution. Their task did not introduce variability in the force trajectories studied, and it does not show evidence of continuous modulation of force.

In contrast, in fNIRS, higher levels of grip forces seem to engage both hemispheres while lower levels mainly engage the contralateral hemisphere (Shibuya et al., 2008; Derosière et al., 2014; Shibuya et al., 2016). This also contrasts with the more frequent lateralisation found in finger pinch tasks (Suto et al., 2002; Nambu et al., 2009; Iso et al., 2016) compared to power-grips (Shibuya et al., 2008; Derosière et al., 2014; Shibuya et al., 2016; Wriessnegger et al., 2017). As a consequence, the loss of spatial resolution for higher forces and power-grips might hinder the decoding of a wider variety of forces with fNIRS. Nonetheless, as we saw in Section 2.2.1, this might also be a consequence of the lack of understanding about the relationship of the HR with force.

Muscle precision In terms of force generated with precision, further evidence is introduced by EEG studies. Iturrate et al. (2018) found that precision and power grips evoked different neural activities appearing bilaterally 100 ms before the onset of the task. In particular, a power-grip elicited higher voltage changes in the posterior contralateral area at the beginning of the task compared to the precision grip. In contrast, the precision task showed higher power in alpha and beta over the ipsilateral centroparietal channels and engaged a frontoparietal ipsilateral network more strongly than the power grip. This higher ipsilateral alpha power has been suggested to reflect inhibition of the parietal regions (Klimesch et al., 2007). It has been found when subjects inhibit movements (Sauseng et al., 2013), leading to lower cortical excitability (Sauseng et al., 2009). The ipsilateral cortex's role seems to be more critical just before the grasping is secured (Iturrate et al., 2018), suggesting that the ipsilateral cortex finely tunes the motor command originated in the contralateral cortex via transcallosal mechanisms in precision tasks (Davare et al., 2006). This effect further increases with the muscle selectivity of a task (Verstynen et al., 2005).

Interim summary In this section, we have seen how different aspects of a task can influence force decoding. While neuronal ensembles can perform well in the decoding of forces for various tasks, ECoG and non-invasive techniques seem to perform better when force is accompanied by movement. We have also seen that force decoding depends on the brain representation of the engaged muscles and the timing of their activations during the task. Ultimately, it is clear that we should design force decoding tasks with all these aspects in mind. Furthermore, the task's characteristics will shape the conclusions we can gather about force decoding capabilities and the strength of our claims. Altogether, isometric tasks seem more relevant to ensure that force –and not movements co-varying with force– are being decoded. Isometric tasks can be even more relevant in non-invasive approaches where it is less clear how the force modulates brain signals. For the same reason, the task should introduce enough variability so that the modulation, if it exists, can be revealed. Another way of ensuring that force, and no other variable, is being decoded is to simultaneously generate independent force profiles with different muscles and see the extent of their reconstruction from the same brain signals. However, in the queried literature, non-invasive procedures never attempted this.

2.3.6 Generalisation in BCI decoders

Another important aspect in BCI decoding is the generalisation to different task conditions. If a decoder is trained in a set of force generation conditions and it still works in a different set of conditions, then we can say it decodes specific information about force control and not specific information about the task. Another way of ensuring this, is using tasks that introduce enough variability in the force profiles so that the decoder learns to explain the greater amount of it in a timed manner. A BCI decoder can also generalise across subjects if decode accurately in unseen subjects during training. Respectively, each case offers evidence that the decoder is relying on brain features that are consistent across tasks or that shared across subjects. Thus, it would show that the decoding is exploiting basic neural function signatures and not spurious characteristics of the data.

As in the previous section, the strong correlation between forces and velocities and accelerations during non-isometric tasks makes it difficult to assess if the generalisation is driven by the decoding of kinematic variables, kinetic variables, or both (Fagg et al., 2009; Gupta and Ashe, 2009). Indeed, when movements are allowed, the generalisation to novel tasks improves when both forces and kinematics are used for the decoding (Gupta and Ashe, 2009) but this also points at the tasks not depending very much on the force. Cherian et al. (2011) focused specifically in generalisation. They found that

under different dynamical and postural conditions, M1 signals decoded EMG signals more robustly than kinematic variables, for which the decoding precision halved. While this could suggest that M1 encodes muscle-like characteristics of force, they also clarify that their micro-electrode array mostly captured task unspecific neurons, probably missing most of the neurons directly connected to muscles which could have decoded kinematic variables better. Chhatbar and Francis (2013) further shows that the decoding of forces allows for better adaptation to different dynamics even in the absence of somatosensory feedback. The long-standing controversy as to whether M1 encodes the kinetics or kinematics variables of movement remains thus unsolved (Fagg et al., 2009) with more recent work showing that M1 and S1 can encode both parameters or their combination at different levels (Branco et al., 2019b; Cherian et al., 2011; Chen et al., 2014; Ganzer et al., 2020).

Recent work also points at non-linearities providing better generalisation in the invasive and non-invasive case. In invasive BCI, Naufel et al. (2019) found that adding a non-linearity to the decoder improved the generalisation of the decoder whether or not force was accompanied of movement and independently of the level of force applied. In the non-invasive case, Paek et al. (2019) suggest that the difficulty in decoding isometric force trajectories from EEG could be due to the different levels of force applied. They found that trials calling for lower forces delivered larger decoded forces, suggesting a non-linear relationship between the ERDs and the generated level of force.

Further in EEG, Hayashi et al. (2019) provided limited evidence that it is possible to decode two different discrete levels of isometric contraction. As we advanced earlier, the limitation in Hayashi et al. (2019) arises from the task and the method they use to prove the force dependent modulation of EEG. On the one hand, the task does not offer any kind of generalisation evidence as they always ask the participants to generate the same profile of contraction without variations in contraction intensities or timings. On the other hand, they train two linear decoders to either distinguish relaxation from low level contraction, or low level contraction from high level contraction. However, Hayashi et al. (2019) do not show evidence of these decoders working simultaneously. At best they show they can distinguish force generation events at different force levels but not EEG modulation with force. Nonetheless, they do offer strong evidence of different topographical distributions of brain activity corresponding to each decoder, suggesting that such differences in topographical distributions associated with different force levels could not be captured by a single linear static model.

Congruent with Paek et al. (2019), Shah et al. (2018) also shows that low level forces are more difficult to decode from invasive STN signals using a linear model. This trend seems to be the opposite for fNIRS

for which stronger forces usually lead to more variable and spatially difficult to resolve hemodynamic responses than low level forces (Derosière et al., 2014; Wriessnegger et al., 2017).

As we can see, EEG and fNIRS convey information about force in potentially complementary ways. However, the combination of fNIRS and EEG, to non-invasively decode force generation and its characteristics using non-linear models, remains unexplored.

2.3.7 What methods have been used to decode force?

The study of force generalisation shows evidence of a non-linear relationship between the brain and the muscles during force generation (Ethier et al., 2016; Paek et al., 2019; Hayashi et al., 2019; Naufel et al., 2019). Nonetheless, linear models are more popular to decode force continuously (Carmena et al., 2003; Gupta and Ashe, 2009; Fagg et al., 2009; Nambu et al., 2009; Cherian et al., 2011; Aggarwal et al., 2013; Chhatbar and Francis, 2013; Chen et al., 2014; Fu et al., 2016; Paek et al., 2019; Beni et al., 2020; Foodeh et al., 2020) due to their easier implementation, low computational weight and lower overfitting (see Table 2.2 for a summary of force decoding implementations). In particular first order Wiener filters or GLMs (Carmena et al., 2003; Fagg et al., 2009; Cherian et al., 2011; Chhatbar and Francis, 2013; Chen et al., 2014; Paek et al., 2019; Ganzer et al., 2020), or Kalman filters (Gupta and Ashe, 2009; Aggarwal et al., 2013) are more frequently used used. On a side note, in the classification of force discrete characteristics, linear models such as Linear Discriminant Analysis (LDA) (Pistohl et al., 2013; Aggarwal et al., 2013; Hansen et al., 2014) or other methods, such as Support Vector Machine (SVM) (Hao et al., 2014; Ganzer et al., 2020), are also more popular.

In the case of BCI decoders using ME arrays, the binned firing rate is more frequently used as input feature for the decoder since it reduces the sparse nature of the spiking activity (Carmena et al., 2003; Gupta and Ashe, 2009; Fagg et al., 2009; Cherian et al., 2011; Chhatbar and Francis, 2013) (Tables 2.2 and 2.3). Other approaches use the near field potential recorded by the ME array in addition to the spiking activity. In the ECoG and EEG cases, spectral features, such as power ad phase, distributed in characteristic frequency bands are used as input features (Chen et al., 2014; Flint et al., 2016; Fu et al., 2016; Paek et al., 2019) (Table 2.3). In lower frequencies ($< 3\,\text{Hz}$), the time course of the electrical voltage, known as Motor Local Potential (MLP) (Fu et al., 2016), is also used and with the delta band contribute more significantly to the non-invasive decoding (Paek et al., 2019). In fNIRS there are less well defined features and most approaches use the amplitude time course of the HbO signals (Nambu

et al., 2009) or HbR (Hansen et al., 2014). Other works use instead the difference between HbO and HbR concentrations (Yin et al., 2015; Fu et al., 2016).

In both cases, the invasive and the non-invasive, one of the most relevant design variables is the memory depth of the features, i.e. the amount of history of brain signals (from t to $t-\tau$) used to decode the force at time t. Since electrical measures (ME arrays, ECoG, DBS electrodes and EEG) are faster, they normally have shorter history depths than the slower fNIRS features. In particular, in between 0.1 to 1.3 s of history are used in electrical measures while hemodynamic fNIRS features often span between 15 to 21 s (Table 2.2). (Nambu et al., 2009) represents an extreme case in which not only history but also 4 s of future HbO signal is introduced to decode the current force, resulting in a non-causal decoder.

Sometimes information about the task is introduced and exploited using Kalman filters (Gupta and Ashe, 2009; Aggarwal et al., 2013). Kalman filters allow the introduction of some statistical task regularities as part of the decoding. While the introduction of task information does not consistently improve the decoding using Kalman filters (Shah et al., 2018; Beni et al., 2020), we remark that including this kind of information can flaw an experiment when the task is very regular. Namely, if the task is very repetitive or has a high auto-correlation, introducing information about the task can bias the decoding to depend on the task and not the brain signals. Namely, a model would find it easier to memorise or auto-encode such an invariable task. In this case, after the model is trained, an analysis should be carried out to clarify the weight that the brain signals have in the decoding. Namely, we would need to understand how much of the performance is only due to the task's auto-encoding and how much is due to the brain signals. To avoid this and other kinds of overfitting, some linear models introduce sparsity in their weights (Fagg et al., 2009; Nambu et al., 2009; Pistohl et al., 2013; Chen et al., 2014). In particular, Foodeh et al. (2020) shows the benefit of introducing sparsity in the linear case.

Since human motor behaviour is very rich in control modes, we are ultimately interested in finding the general rules and not task-specific ones. Indeed, we are constantly switching motor behaviours, for instance, from sitting to standing, from standing to walking and then to running. Each motor behaviour might require different modes of brain control that enable these behavioural differences. These changes in behaviour also represent discontinuities that a single linear model would be unable to capture. Despite the early predominance of linear models, there is a trend in literature gearing towards exploring non-linear methods to decode force, starting approximately in the year 2014 (Table 2.2).

Intuitively, a model for each type of motor control could better capture the relevant brain features for each control type. While a linear model would struggle to capture discontinuous modes of control, separate linear models could "locally" encode each different control mode, which is the idea behind piece-wise linear models (Sachs et al., 2015; Hayashi et al., 2019). Each model for each behaviour is linear, but the transition between them is discontinuous, or potentially, non-linear. In particular, this kind of control requires detecting the transition between models for each control mode. Sachs et al. (2015) use two Wiener models weighted by the likelihood of each control state included in the model, in this case, static posture or reaching. The likelihood of each state controls what model is more relevant at each stage of the decoding. The main advantage was noted in the reach's stability, measured as the ability to stay in the objective during the most prolonged periods. Hayashi et al. (2019) also used a two-stage regression approach to decode two different levels of forces, suggesting that the transition from relaxation to contraction and from one contraction to a higher contraction was not linear. However, Hayashi et al. (2019) did not provide a mechanism to detect when the transition between models happened and depended on this information to work.

Another popular non-linear model is the cascade Wiener Filter (Flint et al., 2014; Ethier et al., 2016; Shah et al., 2017; Flint et al., 2016; Shah et al., 2018). Cascade Wiener filtersindexCascade Wiener filter expand the Wiener filter using polynomials on the regression variables, usually to the third-order (Flint et al., 2014; Shah et al., 2017). Other methods include an additional sigmoid non-linearity in the output of the Wiener cascade (Ethier et al., 2016).

Finally, neural networks have also been used in the decoding of force, but their use is limited to invasive BCIs in monkeys (Naufel et al., 2019). In particular, Naufel et al. (2019) uses an Long-Short Term Memory Network (LSTM) and remarks on the effect that the task's variance structure has in the training of these models. Another study using a recurrent model for the study of monkey kinematics also remarked the importance of taking into account the covariance structure of the task (Sussillo et al., 2012). In both cases, the loss used during training is scaled using the task's variance structure to improve the learning. In any case, Naufel et al. (2019) was the only one of the two studies targeting force generation while Sussillo et al. (2012) targeted kinematics.

In the non-invasive case, neural networks have also been applied to decode discrete characteristics of motor control (Yin et al., 2015; Schirrmeister et al., 2017; Trakoolwilaiwan et al., 2017; Ortega et al., 2018) but, to the best of our knowledge, these have never been applied in the study of force generation. Another shared characteristic across these studies is that, except (Schirrmeister et al.,

2017) in the classification of hand kinematics with EEG, the previous neural networks methods do not offer interpretations of the trained models, and therefore do not advance the field of neuroscience with the application of more complex non-linear methods. However, DL remains unexplored in the study of human force generation with only one example studying force generation with DL from monkeys invasive signals, reaching an approximate 70% of isometric force reconstruction across tasks (Table 2.3) (Naufel et al., 2019).

2.3.8 Conclusion

Our ability to control force enables us to interact with our environment meaningfully. From the most basic tasks, such as carrying a glass of water or brushing our teeth, to the most exquisite ones, like imprinting character in a musical note in a piano or a violin, force enables us to stably and successfully interact with our surrounding. Our hands, in particular, play a central role in all these daily interactions. For people with quadriplegia, among any other motor function, hand function recovery is the most critical priority (Anderson, 2004; French et al., 2010). Furthermore, when hand function is not entirely lost, for example, after stroke, its rehabilitation is also prioritised, and force exercises are preferred (Ward et al., 2003).

Invasive approaches have made much progress in understanding how force can be decoded from brain signals, and this has lead to significant breakthroughs like grasp recovery in a person with SCI (Ganzer et al., 2020). However, invasive approaches require expensive and risky surgeries that stand in the way of reaching most of their potential users.

In contrast, non-invasive approaches offer a safe and affordable way of acquiring brain signals. Additionally, the miniaturisation of devices might soon enable greater commodification, reaching most potential users (Casson et al., 2010). Nonetheless, force decoding is less understood in the non-invasive case, and this review identified two opportunities in advancing force decoding in non-invasive BCI. First, despite evidence showing that fNIRS and EEG provide complemental motor control decoding information (Hansen et al., 2014; Fu et al., 2016), to the best of our knowledge, no work has yet explored this multimodal approach in the decoding of hand forces nor the continuous decoding of force. Second, while invasive research in monkeys has started to elucidate the non-linear nature of force control and the benefits introduced by using non-linear methods (Ethier et al., 2016; Shah et al., 2018; Naufel et al., 2019), all non-invasive methods used in force decoding are still linear (Nambu et al.,

2009; Paek et al., 2019).

The use of non-linear methods becomes critical in the multimodal non-invasive case, not only because of the non-linear relationship between the brain signals and the force but also because of the non-linear relationship between the different brain signal modalities. In particular, fNIRS and EEG measure signals of different physical origins and the relationship between the neuronal activity and the hemodynamic response is, itself, non-linear (Buxton et al., 1998; Friston et al., 2003; Daunizeau et al., 2007; Croce et al., 2017). In the case of fNIRS, the same modelling studies show that the Haemodynamic Response (HR) has a non-linear auto-regulation that delivers non-stationary responses that are sensitive to the parameters of the stimuli used in the task (Glover, 1999). Indeed, the non-linearities introduced by the HR might be complicating the understanding between force generation and fNIRS (Shibuya et al., 2008; Derosière et al., 2014; Shibuya et al., 2016; Wriessnegger et al., 2017). In the case of EEG, linear models might be missing important phasic topographical features that might be hindering the modelling of continuous movement variables (Alexander et al., 2013; Paek et al., 2019; Branco et al., 2019a).

Therefore, non-linear decoding methods can provide better tools for analysing fNIRS signals alone and in combination with EEG. In particular, the lack of non-linear methods might be one reason explaining why multimodal approaches have not yet been explored in continuous force decoding (Ahn and Jun, 2017; Hong and Khan, 2017; Hong et al., 2018). DL is a machine learning framework for building non-linear models which have proved useful not only in image, natural language and speech processing but also in the decoding force from invasive brain signals (Naufel et al., 2019) and the decoding of motor variables different from force from non-invasive brain signals (Trakoolwilaiwan et al., 2017; Schirrmeister et al., 2017; Ortega et al., 2018).

This thesis explores and advances the application of DL methods for fNIRS and multimodal decoding of hand force. We provide non-invasive neuroimaging and force decoding tools to evaluate brain function and assist human force control.

Chapter 3

HYGRiP: Multimodal full brain-muscle characterisation of unimanual grip[1]

3.1 Introduction

Brain-Computer Interfaces (BCIs) offer augmentative communication pathways that bypass the peripheral nerves and muscles (Wolpaw et al., 2002). People can use these brain-computer communication pathways to regain agency over their surroundings, acquiring particular relevance when their muscular pathway is partially or entirely damaged.

Since their first demonstration almost 50 years ago (Vidal, 1973, 1977), BCIs have undergone a steady evolution. Invasive BCIs have achieved significant milestones in continuous signal read out from nervous system activity such as speech decoding (Guenther et al., 2009; Bocquelet et al., 2016; Anumanchipalli et al., 2019), robotic or own arm continuous control (Pfurtscheller et al., 2003; Hochberg et al., 2012) and even grip control with sense of touch recovery (Ganzer et al., 2020).

Non-invasive BCIs have also succeeded in the continuous control of trajectories through sequential control of discrete ERDs (Wolpaw and McFarland, 2004; Royer et al., 2010; Meng et al., 2016). However, while force is essential in motor control (Westling and Johansson, 1984; Ostry and Feldman, 2003), the

[1]Published in: (1) Ortega, P., Zhao, T., & Faisal, A. A. (2020). HYGRIP: Full-stack characterisation of neurobehavioural signals (fNIRS, EEG, EMG, force and breathing) during a bimanual grip force control task. *Frontiers in Neuroscience*, 14:919 and (2) Ortega, P., Zhao T. & Faisal A. A. (2018). Directed information relationship between EMG, EEG and fNIRS during isometric hand contractions. In *2018, 28th Annual Conference of the Society for the Neural Control of Movement*. NCM.

reconstruction of trajectory forces from non-invasive measures is still challenging. Even in the offline case, only modest accuracies have been so far reported using EEG (Paek et al., 2019).

Previous attempts at decoding force from non-invasive measures have focused on the classification of discrete force variables using EEG (Jochumsen et al., 2013; Wang et al., 2017). Yin et al. (2015) showed that using multimodal non-invasive measures (i.e. combining EEG and fNIRS together) increased the classification accuracies of different forces features during imagined hand clenching by 1-5% compared to EEG or fNIRS alone. However, most studies still refrain from using multimodal (fNIRS & EEG) systems, in part, due to the lack of algorithms that successfully integrate both measures (Ahn and Jun, 2017).

Previous research in our lab showed that the combination of different measuring modalities (e.g. Thomik et al., 2013; Belić and Faisal, 2015; Xiloyannis et al., 2017) and the use of state-of-the-art machine learning (Walker, 2015; Xiloyannis et al., 2017; Ortega et al., 2018) could help BCI research if data is collected with a machine learning use in mind.

To stimulate the development of advanced multimodal BCI techniques, we here present the ready-to-use multimodal dataset for the study of isometric force production: **HYGRiP**, the Hybrid (fNIRS & EEG) left & right unimanual isometric grip dataset.

Table 3.1 offers a summary of the characteristics of HYGRiP. The companion repository digests the raw data into a format that makes it at a data readiness level suitable for immediate use by machine learning engineers (Lawrence, 2017). Namely, this level means that anyone can use the data without going through a lengthy process of cleanup and reshaping the data. We believe this will facilitate drawing in more data science and machine learning experts to the exciting problem of BCI.

dataset	N	DWT [h]	trials	task	conditions	available at
HYGRiP	14	2.8	328	L/R unimanual	2 (left/right hand)	data/repo

Table 3.1: HYGRiP characteristics. N stands for "number of subjects" and DWT [h] for the "data-worth time in hours," i.e. the total amount of hours that contain exploitable data. HYGRiP consists of a unimanual task where the left/right (L/R) hands are used in different trials.

3.2 Methods and materials

3.2.1 Participants

Fourteen (N = 14, anonymised IDs from 'A' to 'N') healthy, right-handed volunteers participated in this dataset production. Handedness was confirmed by the Edinburgh inventory (Oldfield, 1971) for all participants. None reported a history of neurological, cardio-respiratory or physical disorders. Imperial College Research Ethics Committee approved all procedures, and all participants gave their written informed consent. The experiment complied with the Declaration of Helsinki for human experimentation and national and applicable international data protection rules.

3.2.2 Motor control task

The motor task consists of a left/right hand-grip. Each hand is a different condition in the experiment. The task consisted of 10 consecutive contraction (1.55 s) and relaxation (0.55 s) periods that introduced rapid force changes. Subjects received visual feedback, and we instructed them to exert forces in the 25 – 50% of their Maximum Voluntary Contraction (MVC), following the pace of the 1.55 s contraction/0.55 s relaxation periods. The 25-50% MVC target range acted as a soft margin within which the subjects had to produce a contraction rather than a varying force they had to track. This MVC target range was implemented to reduce the effect of visual feedback during the task, only provided for contractions out of the task range.

Due to the contraction/relaxation velocity, we do not consider each period a single trial but the consecutive 10 periods as a single trial of the task. In this way, we emphasise the velocity of the grip's execution. Note also that the much slower fNIRS signals need longer times to show a response, and using every single period independently of previous ones could hinder the resolution of the response. We also instructed the participants to prioritise the gripping pace rather than accurately matching the visual cues since the latter was too demanding for the contraction levels required.

Subjects sat in front of the computer screen with their arms relaxed and ergonomically hanging down to facilitate the relaxation of all the muscles not engaged in the task (Fig. 3.1A and B). Subjects were instructed to keep this relaxed posture and reminded to maintain it throughout the experiment.

At the start of the experiment, we asked the participants to produce their MVC with each hand.

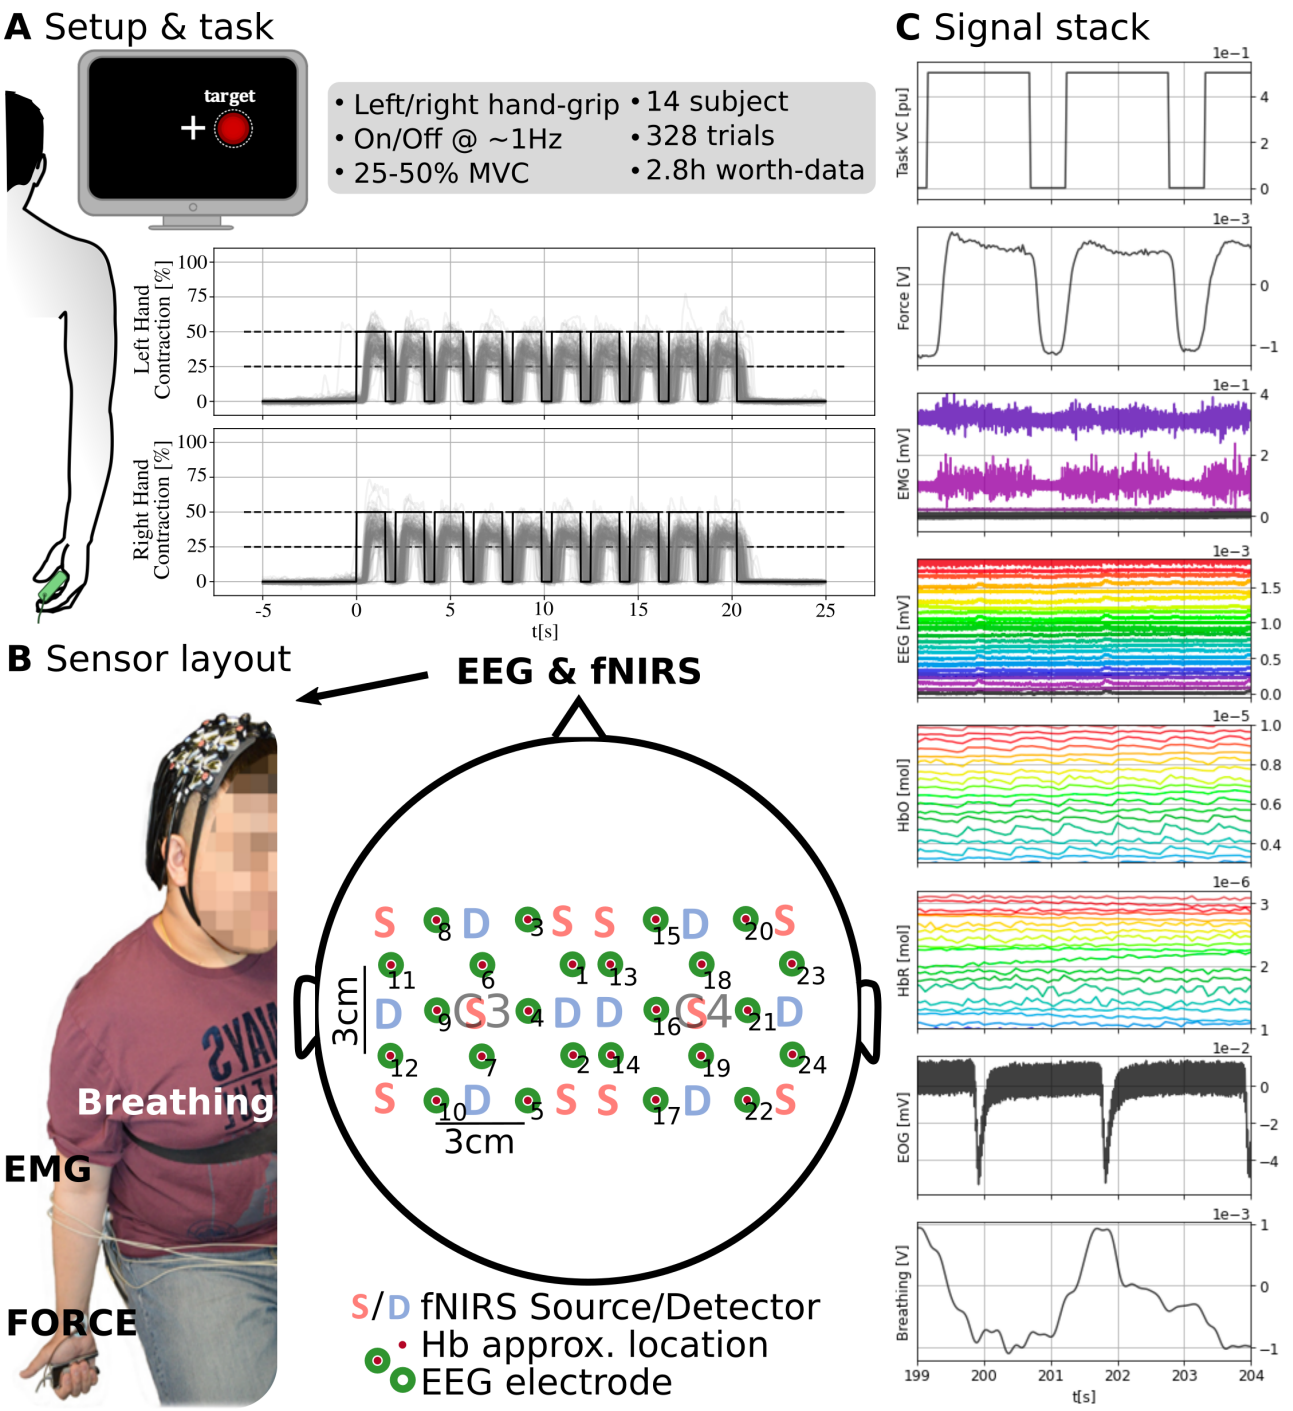


Figure 3.1: **A** (Left) Experimental setup with visual feedback. (Top right) Task and dataset details. (Bottom right) Task execution across subjects time-locked to the 'Go' cue at $t = 0$ s. **B** Sensor layout and EEG and fNIRS sensor arrangement. fNIRS light sources and detectors are placed in a 5×5 symmetrical grid leaving a required 3 cm distance. EEG electrodes, red circles, are placed between each pair of sources and detectors overlaying the fNIRS sensing area so fNIRS and EEG measures are co-located. Each grid is centred around C3 and C4 for the respective hemisphere. The reference sensor was placed in Cz and the ground 3 cm in front of the reference over the mid-line (not shown). **C** A subject's full-stack of neurobehavioural data (5 s, selected channels per modality for readability). (Top-bottom) Force target for the right hand; force produced by subject's right hand; EMG channels from right forearm; EEG channels; Oxygenated Haemoglobin (HbO); Deoxygenated Haemoglobin (HbR); Electrooculography (EOG) with two eye blinks; breathing from chest strap stretch sensor. Units are shown in the bottom x-axis and each corresponding y-axis. Modified with permission from Ortega et al. (2020).

The MVC was used to calibrate the feedback for each hand independently. To obtain reliable MVC estimates, MVC grips were repeated ten times for 1 s with each hand, following paced auditory cues. We computed the MVC as the average of the maximum force across the ten trials. For the experimental task, participants had to generate unimanual hand-grip contractions in the 25-50% MVC range during the 21 s gripping period. The contraction[2] was computed as the ratio Force/MVC. The hand used for the grip varied randomly between trials with equal probability. The experimenter handed the force transducer to the subject within their immediate reach before each trial and removed it after completion.

A computer screen in front of the participants delivered the feedback at an approximate 60 Hz refreshing rate (Fig. 3.1A). A cross indicating the centre of the screen was used as a fixation point during the experiment, and we instructed subjects to stare at it to avoid eye movements. A white discontinuous circumference was used as a visual target and appeared at either side of the cross, ipsilateral to the hand used in each trial. This visual target matched the 25-50% of the MVC of the subject for that hand and paced the contraction-relaxation periods. When the visual target appeared (shown for 1.55 s), the instruction was to contract and, when it disappeared (not shown for 0.55 s), the instruction was to relax. A red filled circle indicated the real-time contraction level, and subjects aimed to fill the discontinuous target circle with red. Visual feedback was only coupled to grips outside the desired range of contraction (25 – 50% MVC), making the red circle bigger or smaller than the target. Otherwise, the red circled filled the white discontinuous circumference within the target range of contraction.

Each trial of the task consisted of 10 consecutive contraction (1.55 s) and relaxation (0.55 s) periods (total, 21 s) with one hand (Fig. 3.1A). Subjects executed the dynamic grip with the hand indicated at the beginning of the trial by synthetic voice 2.5 s before the 'Go' signal (used as the origin of time, $t = 0$ s, for each trial). All participants did a balanced amount of left and right-hand trials of at least 10 (max. 13) trials per hand. We limited the number of trials to avoid muscular fatigue since the task demanded relatively high contractions. Left and right-hand conditions were pseudorandomised across trials to avoid anticipation and interference between conditions. The refreshing rate of the visual feedback overlaid with the force produced in real-time. A randomised resting period, uniformly distributed between 15 and 21 seconds, followed each trial to avoid phasic constructive interference of systemic artefacts (e.g. Mayer waves) in the brain responses.

[2]Throughout this thesis we refer to *contraction* and *force* indistinctly. Note that they have a linear relationship: the MVC factor divides the force to obtain contraction, i.e. contraction = force/MVC. From a BCI perspective, it is always possible to calibrate a decoded contraction range into the desired range of forces.

3.2.3 Data collection pipeline & Methods

We recorded signals representing brain activity, motor behaviour and confounds (Fig. 3.1C). The signals capturing brain activity consisted of EEG and fNIRS. Motor behaviour was captured by the force sensor on which the subjects gripped and surface electromyography on both forearms. We also recorded potential confounds interfering (breathing and EOG) with EEG and fNIRS. We simultaneously used three recording devices. We recorded fNIRS with a NIRScout system (NIRx Medizintechnik GmbH, Berlin, Germany). An ActiChamp amplifier (BrainProducts, Berlin, Germany) was used to record EEG, EMG and EOG. A PowerLab 4/25T system (ADInstruments, Castle Hill, Australia) was used to record force and breathing. The computer used to present the task, and visual feedback was also used to send time-stamps to the three other devices at the beginning and end of the recording and at every 'Go' cue. The time-stamps were stored by each device in its time reference and are used to locate the same event's positions across different devices, allowing us to align the measures to the task events. The sampling frequencies (12.5 Hz for fNIRS, and 4 kHz for remaining measures) were selected to have a common divisor facilitating the resampling processes without the need to round up due to inexact divisors.

Brain signals

We non-invasively recorded all brain signals, and developed a custom 3D printed (formlabs Form2, Formlabs Inc., Somerville, MA 02143, USA) holder made of flexible resin (formlabs RS-F2-FLGR-02) to align the fNIRS and EEG sensors to approximately target similar cortical areas (Fig. 3.1B and 3.2). We configured the sensor layout to result in 12 multimodal (EEG and fNIRS) co-located recording locations per hemisphere. These locations were homogeneously spread with a 3 cm separation creating a grid. Each hemispherical grid was centred around the corresponding 10-20 system C3 and C4 location.

fNIRS We used a NIRScout system (NIRx Medizintechnik GmbH, Berlin, Germany) to record fNIRS signals. We used a total of 12 optodes per hemisphere (10 sources and 8 detectors in total) sampling at 12.5 Hz. An optode is a source-detector pair 3 cm apart from each other (allowing light to reach an approximate 1.5 cm depth into the skull). We placed the fNIRS sources and detectors to result in 12 optodes. The sources and sensors were symmetrically laid around C3 and C4 positions according

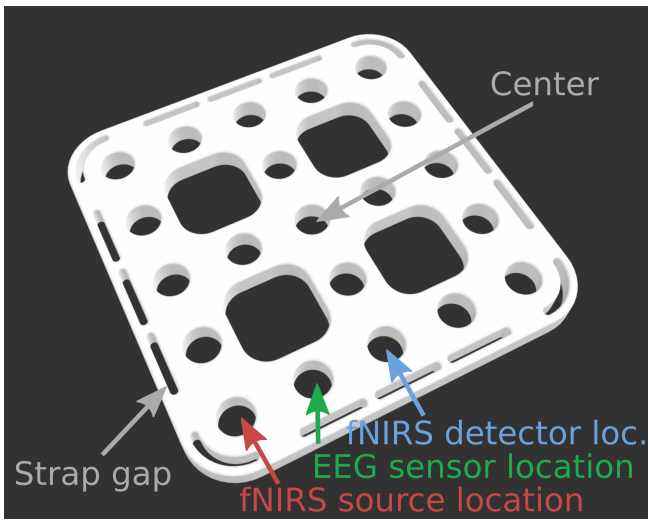


Figure 3.2: Detail of the custom 3D holder design to co-align fNIRS and EEG sensors. We printed two copies of the holder in a flexible resin. Each copy was placed on one hemisphere, and they were kept in position by passing adjustable elastic bands through the gaps. The sensor layout was configured to result in 12 multimodal (EEG and fNIRS) recording locations corresponding to the EEG sensor locations. The EEG sensors and the fNIRS sources and detectors were placed inside the circular holes. The circular holes' locations were homogeneously spread with a 3 cm separation creating a grid. Each hemispherical grid was centred around the corresponding 10-20 system C3 and C4 location. The files required to 3D print the holder are available in the repository for HYGRiP (Table 3.1.)

to the International 10-20 system leaving an inter-optode distance of 3 cm (Fig. 3.1B and 3.2). Two wavelengths (wl_1 = 760 nm, wl_2 = 850 nm) continuous-wave fNIRS was used to obtain optical absorption densities that were transformed to Oxygenated Haemoglobin [HbO] and Deoxygenated Haemoglobin concentrations [HbR] using the modified Beer-Lambert law (Delpy et al., 1988; Cope et al., 1988) (see Eq. 3.2). We also provide the raw optical densities in the dataset.

EEG We used an ActiChamp amplifier (4 kHz, BrainProducts, Berlin, Germany, BrainVision, v1 20 0801) to record EEG. Twelve EEG sensors per hemisphere were placed in between each source-detector fNIRS pair overlaying the region measured by that optode (Fig. 3.1B). We placed the reference electrode in the standard Cz 10-20 position (Nomenclature, 1991; Klem et al., 1999). The signals were low-pass filtered and downsampled to 1 kHz in the dataset.

We used a non-standard sensor arrangement allowed us to align the EEG and fNIRS cortical measures, covering the bilateral sensorimotor cortex (Fig. 3.1B). In our custom sensor holder, every EEG sensor could sit on top of the cortical area targeted by a corresponding fNIRS source-detector pair. Namely, for every fNIRS source-detector pair an EEG electrode was placed in between. Each holder consisted of a 5×5 grid of circular holes whose centres (1.5 cm apart) allowing the required 3 cm separation between fNIRS source-detector pairs and an EEG sensor to be placed between (Fig. 3.2). The positions in the 5×5 grid marked in red in Fig. 3.1B correspond to the physical location of EEG electrodes and the

approximate recording areas of EEG electrodes and fNIRS source-detector pairs.

Physical locations of fNIRS sources and detectors are marked respectively by "S" and "D" in Fig. 3.1B. Twelve (N = 12) recording sites were used per hemisphere due to space limitations, but this was enough to cover the bilateral sensorimotor cortex. Nonetheless, this was sufficient to cover the sensorimotor cortical areas of interest. A holder was placed on each hemisphere and held in position using elastic straps passing through the external gaps (Fig. 3.1B and 3.2). The disposition of every sensing point is symmetrical to the scalp mid-line, and each grid is centred in the corresponding C3 or C4 site depending on the hemisphere side (left or right, respectively) so that the central hole overlaid the respective central position. The C3 and C4 positions were located per subject following the 10-20 standard (Nomenclature, 1991; Klem et al., 1999), i.e. at a 20% of the distance between the pre-auricular points passing over the top of the head from the mid-line. The spherical coordinates of the standard positions are provided in the dataset.

Motor signals

We recorded the grip force and muscular electrical activity to represent the motor behaviour during the task.

EMG We recorded bilateral bipolar surface EMG over the longitudinal axis (+ distal and − proximal) of the muscle belly of the *flexor digitorum superficialis* (4 kHz, on the Aux channels of the BrainVision ActiChamp) placed in the anterior and posterior forearm faces (Fig. 3.1B). Before electrode placement, the skin was cleaned with abrasive pads and alcohol to eliminate dead skin cells and fat impact on electrical recording quality. EMG signals were down-pass filtered and downsampled to 1 kHz in the dataset.

Contraction/Force The dynamic gripping task was conducted using a continuously recorded grip force transducer (PowerLab 4/25T, ADInstruments, Castle Hill, Australia) sampling at 1 kHz. The signal was also used to provide real-time visual display feedback to the subject, and the target force level they were asked to produce (Fig. 3.1A). The MVC force for each subject's hand was measured and computed using the same transducer. The recorded force signals (Volts) were low-pass filtered and downsampled to 50 Hz in the published dataset.

Recording of potential confounds (EOG & Breathing) To complete the picture provided by the data, we recorded potential confounds in the brain and motor behaviour signals of interest. The pulse and breathing rate have an impact at the body level on the concentration of haemoglobin. Therefore, these systemic changes can impact the brain and scalp levels of haemoglobin concentration (Kirilina et al., 2012). While pulse is easily removed in the fNIRS analysis band ($0.01 - 0.25$ Hz), the breathing rate can overlap with it (Pinti et al., 2019). Sources of muscular electrical activity can leak to the EEG sensors and include ocular muscles and skeletal muscles. Thus, we consider electrooculography (EOG), which also carries information on blinks, and breathing as pure confounds. However, EMG might also leak into the EEG sensors and can carry confound information at the same time it provides behavioural information.

Electrooculography (EOG) Bipolar EOG was recorded on the vertical axis (top +, bottom −) of the right eye for all subjects (4 kHz, BrainVision ActiChamp, BrainProducts GmbH, Germany). The signal was low-pass filtered and downsampled to 1 kHz before being included in the dataset. We note that the EOG of participant 'I' is absent. Nonetheless, we included this participant as the impact of the EOG in the recording locations can be less severe than for frontal recording sites and can be corrected using techniques like Independent Component Analysis (ICA) Onton and Makeig (2006).

Breathing We captured breathing as the chest diametrical changes during inspiration and expiration. We used a variable resistor placed inside an elastic strap adjusted around the chest at the level of the Xiphoid process (Fig. 3.1B) to record the expansion and contraction of the thoracic cage (1 kHz, PowerLab 4/25T, ADInstruments, Castle Hill, Australia). The signal was low-pass filtered and downsampled to 50 Hz before being included in the dataset.

3.2.4 Dataset overview

The dataset, provided as a single Hierarchical-Data File (HDF), has undergone very little processing to avoid biasing future analyses. We make the raw data available here and provide companion code that preprocesses the raw data into a readily usable dataset (Data Readiness Level C).

Preprocessing comprises down-sampling to reduce storage space and the formatting of data, recorded events and other meta-data from different devices so that all data followed the same format regardless of their device origin. Further preprocessing can be directly applied using the python packages provided

(Table 3.1), making the data readily available to exploit in popular python machine learning packages such as *pytorch* and *tensorflow*.

Companion code The provided packages only depend on the public python packages: `h5py`, `numpy`, `scipy`, `scikit` and `matplotlib` which need to be installed. The `presentation` notebook contains a thorough explanation and examples of how the tools in `utils` can be used to process the data and depends on `jupyter`. A conda `environment.yml` file is provided with all dependencies to facilitate installation. All together makes the data readily available to exploit, i.e. Data Readiness Level C (Lawrence, 2017).

Dataset organisation The dataset file structure follows a tree-like organisation in three levels. In the first and third levels, the dataset contains meta-data in *string* format that can be accessed via the level's attributes. The first level is the dataset itself and the shared attributes across subjects measures, e.g. sampling frequencies and units, and other information such as the channel grid disposition and a template of multimodal sensor spherical coordinates over the scalp. In the second level, the dataset is organised in one group per subject indexed by their anonymised ID (i.e. fourteen groups with keys `A` to `N`) and contain no attributes. In the third level, each group contains a subgroup for each measure (e.g. keys `frc` for force and `eeg` for EEG) containing the data in *numeric* format and an attribute called `events`, containing the times at which a time-stamp was received during the recording (e.g. "relax", "left-hand", "right-hand"). The events also include the "begin" and "end" time-stamps indicating the beginning and end of the recording session, also *numeric*. In particular, the `events` group contains a second *numeric* attribute `MVC` containing the maximum voluntary contraction value for each hand. For example, for subject "K":

```
Subject K:
    Measure K/brt:   events
    Measure K/eeg:   events
    Measure K/emg:   events
    Measure K/eog:   events
    Measure K/frc:   MVC   events
    Measure K/oi1:   events
    Measure K/oi2:   events
```

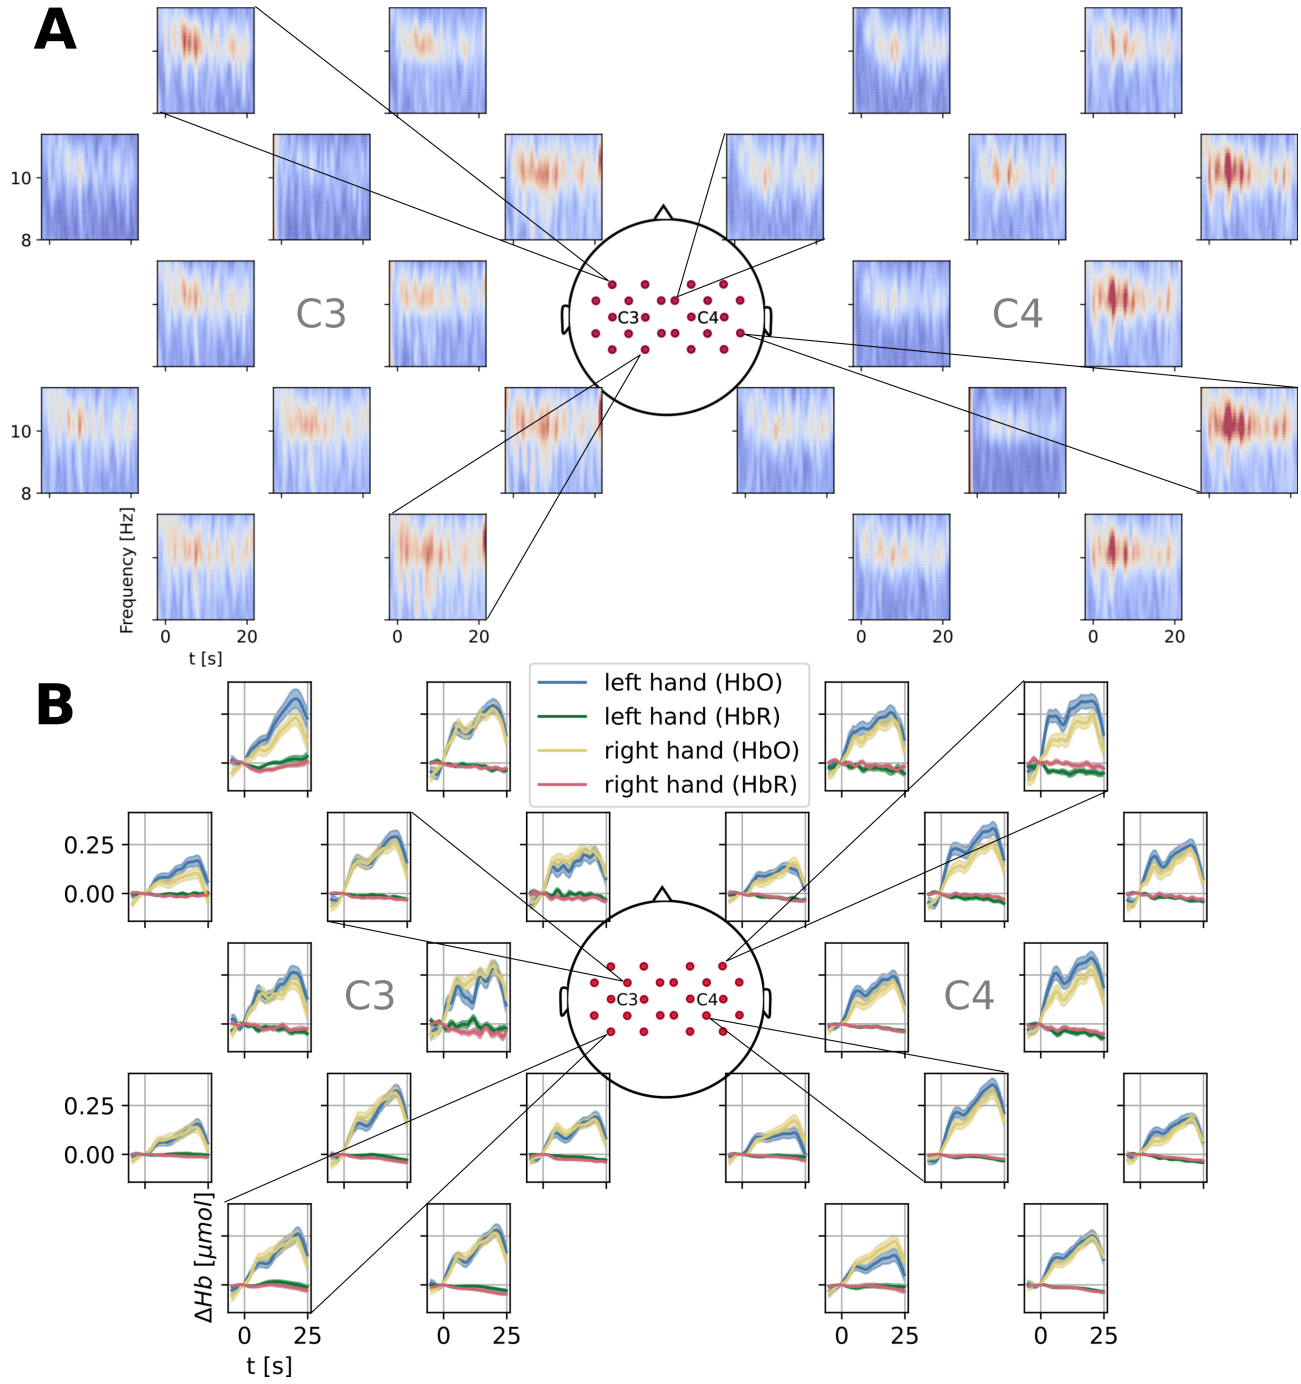


Figure 3.3: **A** Cross-subject mu-band spectrogram averaged across right-hand trials time-locked to the 'Go' cue ($t = 0\,\mathrm{s}$) of each trial. Each trial consists of 10 contraction/relaxation periods (21 s of activity starting at the 'Go' cue). We can observe periodic desynchronisations as alternate red/blue bumps. **B** Cross-subject average of HbO and HbR changes across trials HbO increases during the task, and HbR decreases at a smaller scale. Both start returning to baseline after the 21 s of activity. Units in the bottom-left corner.

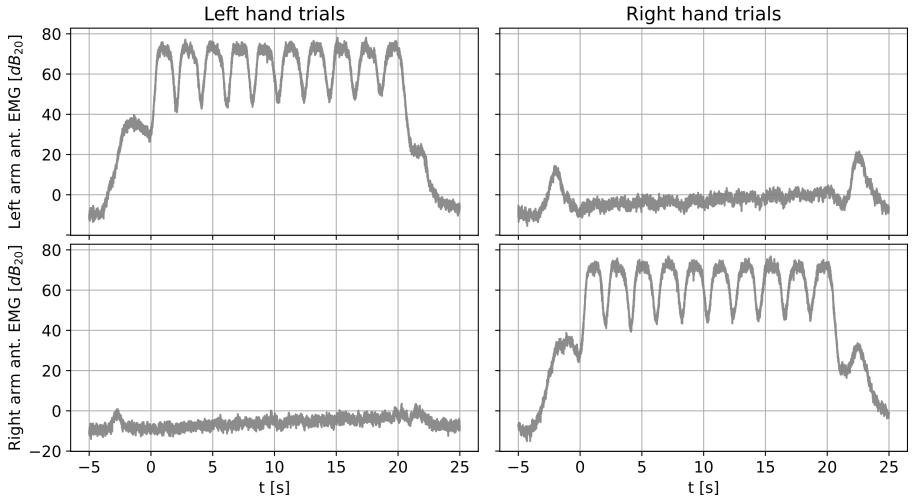


Figure 3.4: Averaged EMG spectral density across trials as in Fig. 3.3A and B, showing similar power density for the active muscles controlling the hands. The passive hands showed higher bumps of activity at the beginning and end of the task when subjects were allowed to relax.

3.3 Data validation

3.3.1 Data validation preprocessing pipeline

The following preprocessing was applied to the raw signals in the dataset to obtain the analysis in Figures 3.3 to 3.5. After preprocessing all signals, we extracted epochs from 5 s before the "Go" instruction to 25 s after.

fNIRS The optical intensity, $\hat{\mathrm{I}}_{ij}^{\lambda}$, for each wavelength, λ, was low-pass filtered below 0.25 Hz with a 7th order elliptical filter. Changes in optical densities per wavelength, $\Delta\mathrm{OD}_{ij}^{\lambda}(t)$, were obtained using,

$$\Delta\mathrm{OD}_{ij}^{\lambda}(t) = -\log\left(\hat{\mathrm{I}}_{ij}^{\lambda}(t) \;/\; \bar{\mathrm{I}}_{ij}^{\lambda}\right) \tag{3.1}$$

with i and j the indices of valid source-detector pairs respectively, t the time and $\bar{\mathrm{I}}_{ij}^{\lambda}$ the average of the optical intensity 1 s prior to the "Go" instruction. Oxygenated and deoxygenated haemoglobin concentration changes, ΔHbO and ΔHbR respectively, were computed solving the modified Beer-Lambert law (Cope et al., 1988; Delpy et al., 1988),

$$\Delta\mathrm{OD}_{ij}^{\lambda} = \mathrm{L}_{ij}^{\lambda}\mathrm{DPF}^{\lambda}(\epsilon_{\mathrm{HbR}}^{\lambda}\Delta\mathrm{HbR} + \epsilon_{\mathrm{HbO}}^{\lambda}\Delta\mathrm{HbO}) \tag{3.2}$$

with DPF^{λ}, the dimensionless *Differential Path-length Factor* accounting for the reduction in intensity due to scattering tissues ($\mathrm{DPF}^{760} = 5.98$ and $\mathrm{DPF}^{850} = 7.54$); $\epsilon_{\mathrm{Hb}}^{\lambda}$, the *molar extinction coefficient* for

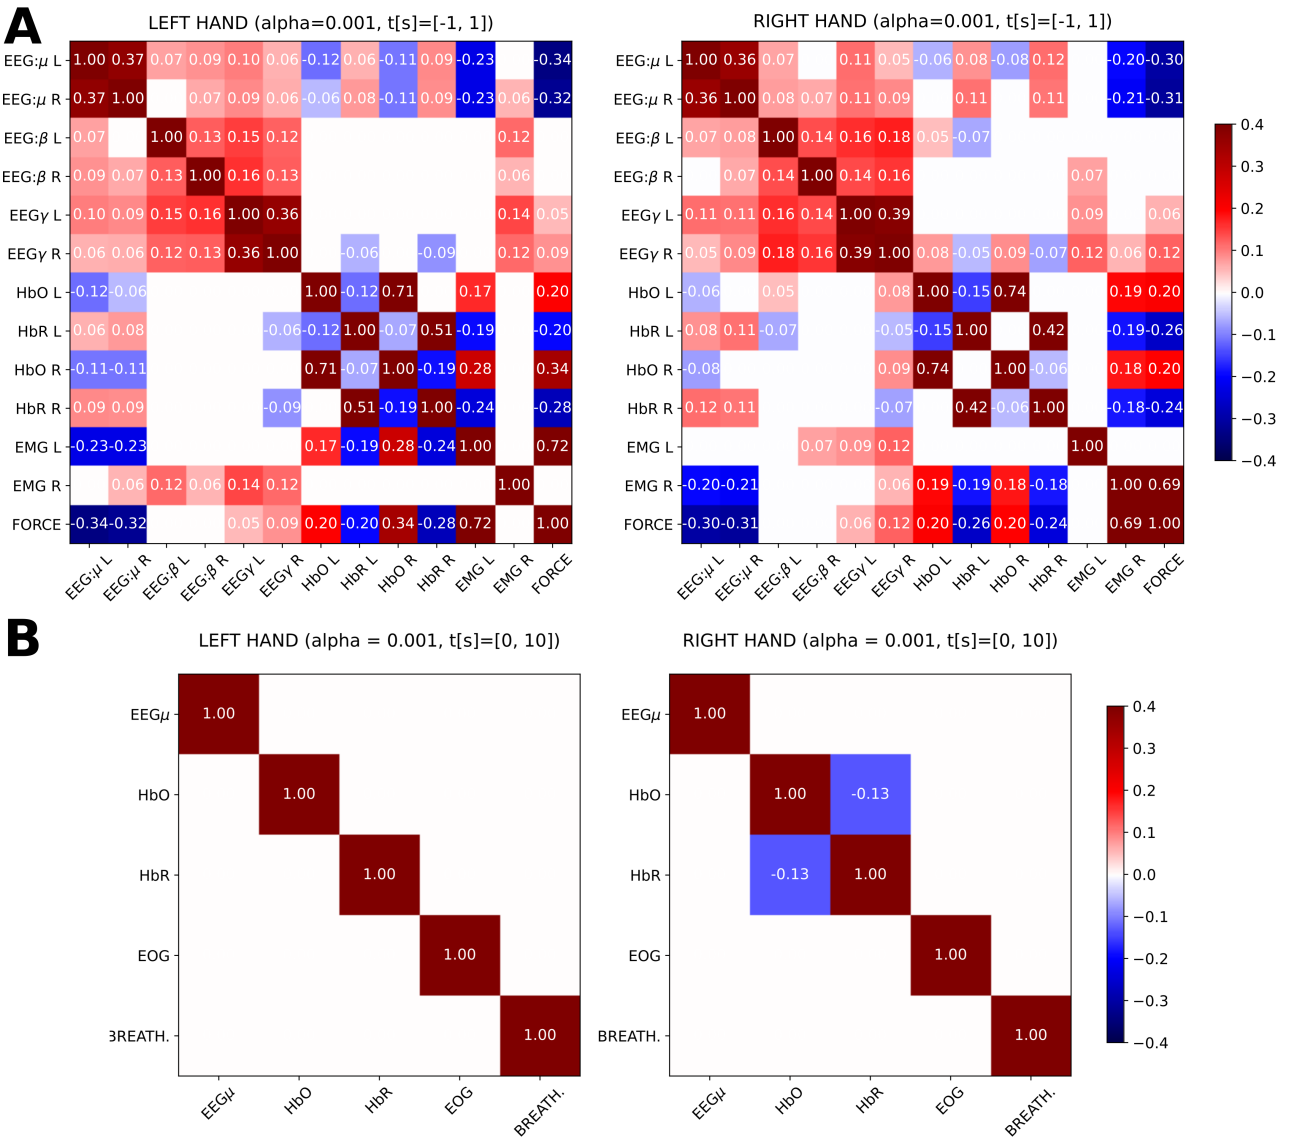


Figure 3.5: **A** Brain and behaviour correlation matrices for the left and the right hand computed on the force onset (−1 s to 1 s around the 'Go' cue). Only significant ($p < 0.001$) correlations are included. It shows a contralateral change in correlation values between HbO and HbR, and force and the time-locked mu desynchronisation event represented as the negative correlation between force and the EEG mu power band. **B** Confounds correlation matrix showing only significant ($p < 0.001$) correlations computed along the task ($t = [-1, 20]$ s around the 'Go' cue) for the right hand condition. Correlations between brain signals and confounds are non-significant ($\alpha = 0.001$).

each haemoglobin and wavelength in $\text{mol}^{-1}\text{cm}^{-1}$ accounting for the absorption of light ($\epsilon^{760}_{\text{HbO}} = 1486.6$, $\epsilon^{760}_{\text{HbR}} = 3843.7$, $\epsilon^{850}_{\text{HbO}} = 2526.4$ and $\epsilon^{850}_{\text{HbR}} = 1798.6$); and L the source-detector distance in cm (L = 3 cm).

After this preprocessing the average and the standard error of the mean across subjects for each hand condition and haemoglobin type were plotted in the corresponding position of the 2D layout (Fig. 3.3B). We can observe an increase and a decrease of HbO with the task onset (t = 0 s) and end (t = 21 s), respectively. We can also identify several peaks in the average response which might be a result of the on-off dynamics of the task which might introduce small variations on the global trend of Hb variations similar to those in the simulations of Fig. 2.6). Opposite changes can be observed for HbR at a smaller scale.

EEG We first downsampled EEG to 250 Hz (with anti-aliasing low-pass filtering), and then applied Notch filters at the mains (50 Hz) and fNIRS (12.5 Hz) frequencies and their harmonics. EEG was finally high-pass filtered above 1 Hz using a 5th order Butterworth filter. ICA was used in two stages to automatically remove components correlated first with EOG and second with EMG. For EOG, both the EOG and EEG were downsampled to 25 Hz. We automatically rejected a maximum of 1 independent component correlated above 0.3 (in absolute values) with the EOG. For EMG, both the EMG and EEG were downsampled to 125 Hz. The rejection of components was more strict to ensure that EMG was not contaminating the data. One component was rejected whenever its correlation magnitude with any recorded EMG channels was greater than 10^{-4}.

Figure 3.3A presents the averaged spectrogram across subjects for the right-hand condition (the left-hand condition can be found in the notebook) for the mu band (8 – 13 Hz). Interestingly, the on-off nature of the task might be introducing periodic variations of power in the mu band (Fig. 3.3A) due to Event Related Desynchronisations (ERDs) (Pfurtscheller and Lopes Da Silva, 1999; Pfurtscheller et al., 2006).

EMG The EMG was also first downsampled to 250 Hz (with anti-aliasing low-pass filtering), then high-pass filtered with a 17th order Butterworth filter of above 110 Hz. To generate Fig. 3.4 we computed the Hilbert envelope of the signal (Cizek, 1970) and used it to obtain decibels (dB) of power density referred to the mean power of the signal during the epoch.

We finally averaged these power densities across subjects per hand condition and arm location of the electrodes. The active EMG (i.e. those corresponding to the arm used during each hand condition)

carry a similar amount of power density for each hand condition. The passive electrodes have a much flatter amount of power density during the task. In contrast, they have clear peaks at the beginning and the end of the trials when the subjects were allowed to relax.

Force The force signal was band-pass filtered between 10^{-4} – 9 Hz (2nd order elliptical filter). Once we extracted the epochs, the force signal was again high-pass filtered above 10^{-3} Hz to remove any remaining offset. These low high-pass frequencies were selected to preserve the squared shape of the forces, very rich in low frequencies. Once the offset was removed, hand contraction values were obtained by dividing the resulting forces by the maximum voluntary contraction force recorded at the beginning of the trial.

Fig. 3.1A shows the gathered trials for all subjects per hand condition. Subjects mainly engaged with the task in timing and contraction values, with the left condition presenting slightly more overshoots. Although the task is conceptually simple, the provision of only partial visual feedback and its fast on-off nature contributed to higher variability in the behaviour within the desired levels of contraction. We consider these aspects of being more representative of natural force applications where feedback is mostly proprioceptive, and changes in force can be fast and span a broader range than discrete target levels.

Other signals (EOG & Breathing) EOG was downsampled (with prior anti-aliasing filtering) to 50 Hz. They were then filtered using an 8th order high-pass filter above 1 Hz. Breathing was also low-pass filtered with a 6th order Butterworth filter below 0.25 Hz.

3.3.2 Correlations across multi-modal signals and validation

Figure 3.6 illustrates the different time-scales of the signals characterising HYGRiP. We can see in particular how the brain signals, fNIRS and EEG, occupy opposite extremes of the band-width. The generated forces are placed in the middle.

Figure 3.6 also depicts the prominent mu and gamma-bands related to motor control in the EEG. To further inspect how these signals, with very different time-scales and natural origins, are related, we compute their correlation. Hence, after preprocessing, signals were correlated to characterise a “brain

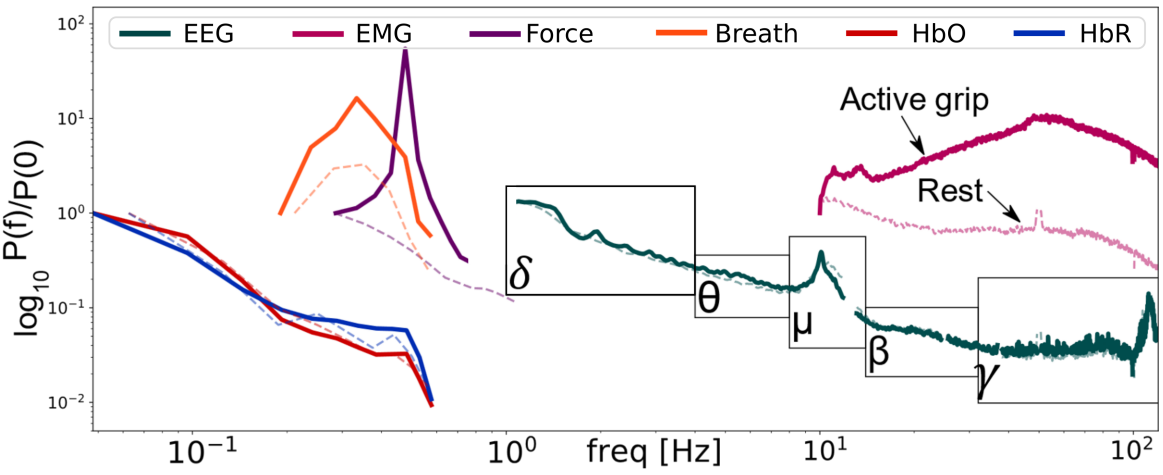


Figure 3.6: Average periodogram of the multimodal signals composing HYGRiP. This figure depicts the broad differences in time scales between the different signals. In particular, fNIRS occupies the lowest frequencies below 1 Hz. The breathing overlaps with fNIRS, and the produced force is in between the fNIRS and the EEG's delta band. We depict the typical EEG-bands as squares with their corresponding Greek character. Note, in particular, the peaks in the mu-band and high gamma-band related to motor control (Chapter 2).

and behaviour" or neurobehavioural correlation including EEG, fNIRS, EMG and force (Fig. 3.5A); and a "confounds" correlation structure including: EEG, fNIRS, EOG and breathing (Fig. 3.5B).

EEG and EMG underwent additional filtering and spectral density computation similarly to that used to produce the EMG plot (Fig. 3.4). For EEG the mu (8 – 12 Hz), beta (12 – 30 Hz) and gamma (60 – 125 Hz) bands spectral densities were extracted. Then, signals were downsampled to the lowest sampling frequency in the dataset, i.e. 12.5 Hz, and cropped between −1 and 1 s ('Go' cue at $t = 0$ s) to focus on the onset of motor activity, in the case of the 'brain and behaviour' correlation. The 'confounds' signals were cropped between −1 and 21 s. All signals were finally normalised.

For each hand condition, the 'brain and behaviour' correlation (Fig. 3.5A) was computed over the appended observations corresponding to the $t = [-1, 1]$ s crops for all combinations of the left and right hemisphere signals. The signals included the mu, beta and gamma EEG power bands, the HbO and the HbR from fNIRS; and the right and left arm EMG and force. For the brain signals, channels 4 and 16 (Fig. 3.1B) were selected as representative of the corresponding hemisphere activity. A similar process was used for the 'confounds' correlation for the $t = [-1, 20]$ s. A significance level $\alpha = 0.001$ is set, and only significant correlation values are shown.

As we can observe, there is a minimal correlation between brain signals and confounds (< 0.03 in absolute value, not shown due to lack of significance $\alpha = 0.001$, Fig. 3.5B), suggesting that the confounds do not interfere with the recorded signals after applying the preprocessing pipeline.

Our data's key observations can be found in the 'brain and behaviour' correlation (Fig. 3.5A). The

first and most obvious observation is the strong correlation ($r \approx 0.7$, $p < 0.001$) between the EMG power envelopes of the active hands and the force, which supports the synchronisation of the different devices used to record these measures.

To analyse the interaction between brain signals (fNIRS and EEG) and task measure (force), we need to bear in mind that the time scales of fNIRS and EEG are very different (seconds vs milliseconds). In the EEG domain, it is known that motor activity onsets are reflected in EEG power features. In our data, we observe significant strong anticorrelation ($r \approx -0.3$, $p < 0.001$) between the EEG power in the mu band and the force, which indicates that passing from a resting state (high mu power) to motor activity (low mu power) is adequately captured by these spectral features when we look at a time window from −1 s to 1 s around the "Go" cue. This decrease in power, also known as *mu event-related desynchronisation* (mu-ERD), is due to the desynchronisation of neuronal activity (Pfurtscheller and Lopes Da Silva, 1999; Pfurtscheller et al., 2006), and shows that the EEG is aligned with the force, helping to strengthen the validity of our dataset.

However, when the correlation is computed focusing on the 20 seconds of the task (from 0 s before the Go cue to 20 s after, not shown), the mu-ERD is not significant ($p > 0.001$). Only the beta band (12 – 30 Hz), known to be synchronised with motor activity (Kristeva-Feige et al., 2002), appears with a low level of positive correlation ($r \approx 0.04 - 0.05$, $p < 0.001$). The lack of mu-ERD and force correlation during the 20 s of continuous contraction/relaxations might be a consequence of the velocity at which the sequential contractions/relaxations were executed. These fast phase transitions in force might not have left enough time for the motor cortex to reach a synchronised equilibrium state before it was desynchronised again. Furthermore, this can also indicate that mu frequencies (8 – 12 Hz) are not fast enough to track this kind of subtle phase changes and justify developing more precise algorithms or feature extractors (Paek et al., 2019).

In the fNIRS domain ±1 s around the "Go" cue, we also observe a characteristic HbO/HbR anticorrelation ($r \approx 0.15$, $p < 0.001$) (Jasdzewski et al., 2003; Huppert et al., 2006) in the structure. Higher correlation magnitudes are associated with the contralateral hemisphere to the hand used. Namely, the HbO/HbR anticorrelation is more robust in the right hemisphere for the left hand, although it is also present in the left hemisphere. This lateralisation suggests that the dominant right hand may engage the left hemisphere more strongly while the non-dominant left hand engages both hemispheres in a more balanced way.

HbO and HbR also show significant correlation with the force ($r \approx 0.2 - 0.3$, $p < 0.001$). In particular, HbR, which is more specific than HbO (Hirth et al., 1997), also shows higher anticorrelation with force for the contralateral hemisphere. This brief analysis aims to validate the dataset and present some of its features.

We will come back to HYGRiP in the following chapter, where we develop a DL method that deepens in the explanation of the relationship between fNIRS and left/right hand grips. To date, the notion prevails that hand grips evoke a global response over the motor cortex, i.e. both hemispheres show the similar HR during force grips (Shibuya et al., 2008; Derosière et al., 2014; Shibuya et al., 2016). However, it is suggested that the relationship between force and HR is more complex than previously thought, and these differences may lie in complex features of the HR (Wriessnegger et al., 2017). In next chapter (Chapter 4), we will show how DL can be used to resolve hand specific HRs activations in the cortex as measured with fNIRS.

3.4 Conclusion

This chapter has presented a thoroughly collected and compiled multimodal brain and behaviour dataset to study the brain representation of force non-invasively. The companion code makes the data easily exploitable by the neuroscience and machine learning communities. In this line, HYGRiP is already published in Ortega et al. (2020). Our aim with HYGRiP is to encourage the community to develop algorithms to understand better the rich temporal relationships between brain signals spanning very different time scales and physical origins with BCI and neuroimaging purposes. We believe that the availability, ease of access and thorough documentation of this complex dataset might inspire other machine learning researchers to join the challenge, ultimately accelerating the development of more sophisticated signal processing and decoding algorithms for multimodal brain signals.

Chapter 4

HemCNN: Improving fNIRS' temporal and spatial resolution using Deep Convolutional Neural Networks during strenuous force[1]

4.1 Introduction

Non-invasive neuroimaging techniques provide an easy way to study the human cortex in a variety of tasks. Significantly, in sensorimotor tasks that reflect on ecologically valid real-world tasks, these methods can help us evaluate, rehabilitate, replace or assist the motor function after brain damage (Calautti and Baron, 2003; Daly and Wolpaw, 2008; Dimyan and Cohen, 2011; Enriquez-Geppert et al., 2013). In this context, two-wavelength functional Near-Infrared Spectroscopy (fNIRS) is a cost-efficient and portable neuroimaging technique that measures the cortical Haemodynamic Responses (HRs) following the oxygen demands of active neurons.

However, in some cases, the HR is not easy to resolve. Several sources of noise (e.g. pulse, breathing

[1]Published in: (1) Ortega, P. & Faisal, A. A. (2021). HemCNN: Deep Learning enables decoding of fNIRS cortical signals in hand grip motor tasks. In *2021 10th International IEEE/EMBS Conference on Neural Engineering*. IEEE. (2) Ortega, P. & Faisal A. A. (2019). Deep learning time-resolves fNIRS cortical activation asymmetries during a force tracking task evoking non-stationary haemodynamic responses. In *2019 49th International Conference of the Society for Neuroscience*, program No. 760.15. SfN. Submitted as journal article to: Ortega, P. & Faisal, A. A. HemCNN: Deep Learning enables decoding of fNIRS cortical signals in hand grip motor tasks. *Neuroimage*.

and Mayer waves) increase the variability of the HR and reduce its spatio-temporal resolution (Kirilina et al., 2012; Tachtsidis and Scholkmann, 2016; Caldwell et al., 2016). Linear techniques, like GLM in SPM, have proven very useful in the reduction of these sources of noise (Friston et al., 1994; Schroeter et al., 2004; Ye et al., 2009). However, these linear techniques require previous knowledge of the relationship between the task under study and the HR. For a given task, these techniques first compute an expected HR and then regress out all the variability in the real fNIRS measurements that does not fit the expected HR. Therefore, these linear techniques are limited to cases where the relationship between the task and the evoked HR is known. Otherwise, the application of these linear methods can result in HR variability being discarded as noise when it should have been explained by neuronal activity (Faisal et al., 2008). This limitation can be particularly relevant in the study of motor control where it has been observed that the speed (Obrig et al., 1997; Kuboyama et al., 2005) and the force (Shibuya et al., 2008; Nambu et al., 2009; Derosière et al., 2014; Wriessnegger et al., 2017) of a motor execution can modulate the HR, yet the exact analytical relationships are unknown.

Apart from the lack of knowledge about the task-HR relationship, HR physiological models (Buxton et al., 1998; Friston et al., 2003; Daunizeau et al., 2007; Croce et al., 2017) are also non-linear in nature and predict that any stimulus driving the system faster than its time constant (i.e. faster than 0.65 Hz) can give rise to non-stationary responses that increase the variability of the HRs (Glover, 1999).

Altogether, these studies suggest that data-driven and non-linear methods can be more suitable for capturing relevant HRs variability from fNIRS during force generation. In particular, in force generation, it has been observed that changes in the task parameters can be one of the reasons explaining why some force-tasks evoke spatially-resolved HRs (finger-pinches, Nambu et al., 2009) while others lose this spatial resolution (hand-grips, Shibuya et al., 2008; Derosière et al., 2014; Wriessnegger et al., 2017) when linear methods are used. In the latter studies, this lack of spatial resolution led to inconclusive differences across brain hemispheres, which is an often sought result during motor control (Penfield and Boldrey, 1937; Serrien et al., 2006). To advance the neuroimaging capabilities of fNIRS in human force control, we set out to explore how DL can be used to improve the decoding and interpretation of non-stationary HRs. In this chapter, we develop hemCNN[2], a DL approach based on CNNs. We use hemCNN to decode spatio-temporal hand-specific cortical activities during force generation. In particular, hemCNN is designed to learn differences across hemispheres by learning to classify a motor task with two classes. Note that the multi-class case is beyond the scope of this work.

[2]HemCNN stands for Hemispherical-CNN since it is designed to learn differences across brain hemispheres.

CNNs are a type of feed-forward neural networks that find optimal convolutional filters to process their input signals. They learn to enhance distinctive features present in the original data, to facilitate, for example, their classification. CNNs have been previously applied to fNIRS by Hiwa et al. (2016), to detect gender based differences, by Trakoolwilaiwan et al. (2017), to classify left and right-hand finger taps, and by Ghonchi et al. (2020) to classify motor imageries and mental arithmetic. However, none of them studied highly variable HRs evoked during force generation nor provided a physiological interpretation of the learned features. Here, we embed research hypotheses in the hemCNN architecture that are tested and interpreted across subjects. Namely, we hypothesise that (1) hemCNN is better than linear methods at processing highly variable HRs evoked during fast hand-grips, that (2) the HbO and HbR fNIRS measures provide a functional signature of cortical activation, and that (3) hand-specific cortical spatio-temporal activity patterns are associated to the left and the right-hand grips.

4.2 Methods

Our aim here is to show how DL can advance neuroimaging during force generation in humans. Therefore, the core methods relate to the design, training and interpretation of hemCNN. Nonetheless, to provide a complete picture, we start analysing some behavioural aspects during force generation and explaining the standard fNIRS processing techniques that we use to benchmark the hemCNN results.

4.2.1 Data

The data used for this chapter comes from the Hybrid (fNIRS & EEG) unimanual (left & right hand) isometric Grip dataset (HYGRiP) dataset[3]. The dataset contains fNIRS, EEG, EMG, force and breathing measures during a fast hand-gripping task for 14 subjects. Here only fNIRS, force and breathing measures are used.

HYGRiP's task consists of periodic unimanual grips at an approximate rate of 1 Hz with contraction levels between 25 – 50% MVC (Fig. 4.3) which are performed with either hand at different times. Two subjects were excluded for having larger maximum levels of contraction for the left hand (58% and 57% MVC respectively) than the maximum indicated in the task (50% MVC), leaving twelve (N = 12)

[3]We describe here the relevant details. For a detailed description of the dataset, refer to Chapter 3.

subjects to analyse. It is expected that the rate of the contractions (~ 1 Hz and above the time constant of the HR, 0.65 s^{-1}) and the high levels of contraction (25 – 50% MVC) of the task introduce variability in the HR. At the same time, this rate and level of contractions represent a less constrained use of force and, in this sense, are more representative of naturalistic force generation.

A trial consists of a ~ 1 Hz paced on-off unimanual force tracking of 21 s duration (Fig. 4.3A). Each trial was followed by a randomised resting period uniformly distributed between 15 and 21 seconds, to avoid phasic constructive interference of any systemic artefacts in the brain signals. The first 10 left-hand trials and the first 10 right-hand trials per subject were used.

For the analyses presented in this chapter, we used the fNIRS, the force, the EMG and the breathing signals from HYGRiP. fNIRS is used to analyse brain activity during the grip control task. EMG and force are used in the behavioural analysis to verify that both hands had similar contraction levels during the task. The breathing signal was used to control for possible hand-use induced systemic artefacts that could contaminate the fNIRS signals.

4.2.2 Behavioural analysis

With this brief behavioural analysis, we aim to verify that subjects perform the task as expected and that both hands have similar levels of contraction using the force and the EMG signals in HYGRiP.

Force Raw force signals from the HYGRiP dataset are band-pass zero-phase filtered from 0.1 Hz to 5 Hz. The MVC value of each hand obtained during calibration is used to compute the contraction values (per unit, p.u.) for each hand,

$$\text{contraction [p.u.]} = \frac{\text{force [V]}}{\text{MVC [V]}} \tag{4.1}$$

Since the measures did not meet normality assumptions, a Mann-Whitney U test was used to check that the left and right-hand contraction levels had similar distributions (see Figure 4.3A right).

Electromyography Raw EMG signals from HYGRiP are downsampled to 500 Hz, notch-filtered (50 line noise and its higher harmonics 100, 150, 200 Hz) and zero-phase high-pass filtered above 60 Hz to reduce low-frequency interference with the task. The EMG power profile is obtained as the

spectrogram's average power in the $60-200\,\text{Hz}$ range at each time window in the spectrogram. The left and right-hand populations of such gathered profiles are compared using a Mann-Whitney U test.

4.2.3 fNIRS traditional analysis

To build a benchmark for hemCNN, we here use conventional fNIRS analysis techniques. For this analysis, the fNIRS optical densities from HYGRiP were downsampled to $2\,\text{Hz}$ and zero-phase band-pass filtered between $0.01-0.3\,\text{Hz}$ to remove drift and pulse artefacts.

Haemoglobin change quantification (ΔHb) fNIRS measurements of Hb concentration do not provide an absolute measure of brain activity. To provide a more specific quantification of brain activity in the conventional approaches, we refer the hemodynamic activity to no-activity periods (e.g. a resting time) that usually occur before a particular behaviour under study starts (Shibuya et al., 2008; Derosière et al., 2014; Wriessnegger et al., 2017), known as baseline correction.

For traditional analyses and standard methods, Hb time-series are baseline corrected using the average of the $5\,\text{s}$ before stimulus onset, $t = 0\,\text{s}$. The averaged values during the last $5\,\text{s}$ of stimuli of channel 7 (left hemisphere) and channel 19 (right hemisphere, Fig. 3.1B) are respectively used as left and right brain activation changes (Fig. 4.4C).

For example, for the left-hand,

$$\Delta\text{Hb}^L = \langle \text{Hb}^7(t) \rangle_{t=[15,\,20]\ \text{s}} - \langle \text{Hb}^7(t) \rangle_{t=[-5,\,0]\ \text{s}} \tag{4.2}$$

where $\langle \text{Hb}^i(t) \rangle_{t=[a,b]}$ denotes the time average between time $t = a$ and $t = b$ of the signal x on channel i. These Hb changes are used to compute the Lateralisation Index (LI) and for classification with conventional methods.

Lateralisation index (LI) The LI is a widespread measure to study hemispherical dominance (Seghier, 2008). Here, LIs are used to compare the hemispherical dominance associated with each hand. $LI > 0$ indicates right hemisphere dominance and $LI < 0$ left hemisphere dominance,

$$LI_{\text{HbO}} = \frac{\Delta\text{HbO}^R - \Delta\text{HbO}^L}{|\Delta\text{HbO}^R| + |\Delta\text{HbO}^L|} \tag{4.3}$$

where the indices L and R correspond to the left and right hemispheric location of the channel for which the LI is being computed. Since ΔHbR are mostly negative, to achieve consistency, i.e. $LI > 0$ indicating the dominance of right hemisphere activity, LI_{HbR} is multiplied by -1. A Mann-Whitney U test is used to find differences between LIs distributions across hands.

4.2.4 hemCNN analysis

HemCNN, a DL approach to advance fNIRS neuroimaging during force control, is the main contribution of this chapter and is presented here. Bold Latin lower case letters denote vectors, $\boldsymbol{x}$, and bold capital letters denote matrices, $\boldsymbol{X}$.

Conventional fNIRS features for classification

The following methods are used to compare and assess the performance of hemCNN and represent widely used methods to process fNIRS signals. HemCNN's performance is measured in terms of its accuracy at predicting the true label l (i.e. the hand used to generate force, $l = 0$ for the left and $l = 1$ for the right) from the brain signals of subjects on which hemCNN was not trained on. We define this test accuracy as,

$$\text{Test Acc.}[\%] = \frac{\sum_{i=0}^{N} \delta(l - \hat{l})}{N} \cdot 100 \tag{4.4}$$

With $\hat{l}$ the predicted label, δ a function valued 1 in l and $\hat{l}$ are equal, or 0 otherwise, and N the number of examples.

We use these conventional methods to extract features to train a simple tree classifier. These classifiers use the same leave-one-subject-out test strategy and the same classification task that we use to train hemCNN. In contrast to hemCNN, all the standard procedures here use a baseline corrected fNIRS input. Namely, fNIRS signals are referred to the average of the 5 s before the go-cue in the task. This baseline correction increases accuracy, though not significantly. We decided to keep the baseline correction for the conventional procedures.

General Linear Model GLM is the current standard in fNIRS analysis (Schroeter et al., 2004; Plichta et al., 2007; Koh et al., 2007; Ye et al., 2009). We use it here to benchmark our hemCNN approach. The GLM design matrix, $\boldsymbol{D} \subseteq \mathbb{R}^{n \times m}$, contains n canonical HRs, e.g. gamma functions (Ye

et al., 2009), of m time samples each. These are multiplied by a set of weights, $\boldsymbol{v} \subseteq \mathbb{R}^n$, to approximate the input signal $\boldsymbol{\Delta Hb} \subseteq \mathbb{R}^m$. Here **Hb** denotes either oxy or deoxy-hemoglobin concentrations time-series for one channel. The Δ increment emphasises that the signals are referred to the baseline before the stimuli start. The resulting model equation is,

$$\boldsymbol{\Delta Hb} = \boldsymbol{D}^T \boldsymbol{v} + \boldsymbol{\epsilon} \tag{4.5}$$

with $\boldsymbol{\epsilon}$ a vector of m samples representing the noise with normal distribution. The design matrix $\boldsymbol{D}$ weighted by $\boldsymbol{v}$ is used to reconstruct the $\boldsymbol{\Delta Hb}$ time-series of each individual channel. To account for HbO and HbR channel relationships and provide a more fair comparison to CNN capabilities of processing multivariate signals, we introduce a variation of GLM. The idea is to capture the first principal components in two hierarchical steps using Principal Component Analysis (PCA). We call this approach GLM with hierarchical principal component analysis or GLM-hPCA.

First, PCA is performed in the paired HbO and HbR features of the same channel, $\boldsymbol{v}$. The first PCs of such interactions are aggregated per hemisphere, and we run PCA again. Finally, the first PC of each hemisphere is kept as a unique feature per hemisphere. In both GLM approaches, five ($n = 5$) canonical gamma functions rising and decaying in chronological order across the 20 s stimulation-period are used as the design matrix for HbO. The opposite sign of the previous design matrix is correspondingly used for HbR. After GLM, each channel of each 20 s trial segment is represented by a $\boldsymbol{v} \in \mathbb{R}^5$. Each element in $\boldsymbol{v}$ represents approximately $\frac{1}{5}$ of the segment's time-series. Therefore, the five features represent a chronologically ordered abstract representation of the time-series, from further in the past to most recent. For the standard GLM approach, the total dimension is 240 (5×48 HbO and HbR channels), and for GLM-hPCA this amount is reduced to 10 (5×2 hemispheres).

ΔHb and LI The measures used in the traditional characterisation of the HR are also used as features for classification. ΔHbO and ΔHbR, each rendering a 1-dimensional feature per channel, are aggregated, resulting in a 48-dimensional feature per example[4] (2×24 HbO or HbR channels). LIs are computed per each pair of symmetric channels leading to a 24-dimensional feature per example.

[4]We use "example" in this chapter to denote a sample from the dataset fed into a classification algorithm. Thus, an example corresponds to an experimental "trial". This is to make a clear distinction between a "time-sample" in the time-series and an "example", also called "sample" in literature, from the dataset.

HemCNN design and training

Convolutional Neural Networks (CNNs) optimise convolutional filters (CF) to extract relevant features for a machine learning task, e.g. classification of examples into categories. Many fNIRS processing steps (e.g. Hb feature description and channel selection) can be implemented into an appropriate CNN architecture that would thereafter optimise filter weights to increase classification accuracy across the different conditions of an experiment. We designed the hemCNN architecture and its training to enhance hemispherical differences between fNIRS signals evoked by different hands.

Specifically, the classifier corresponds to a mapping

$$\boldsymbol{f} : \mathbb{R}^{\mathrm{nCh} \times \mathrm{nSamps}} \longrightarrow \mathbb{R}^2 \tag{4.6}$$

with nCh the number of channels, 48, nSamps the number of time samples, $19\,\mathrm{s} \times 2\,\mathrm{Hz} = 38\,\mathrm{samples}$, and 2 being output dimension which correspond to the two hands. In our terminology each element of the input space, $\boldsymbol{X} \subseteq \mathbb{R}^{\mathrm{nCh} \times \mathrm{nSamps}}$, is an "fNIRS example" for the network. The values of the output units were passed through

$$\mathrm{LogSoftMax}(x_i) = \log\left(\frac{\exp(x_i)}{\sum_j \exp(x_j)}\right) \tag{4.7}$$

so that each unit represented the logarithmic probability of an fNIRS example ($\boldsymbol{X} \subseteq \mathbb{R}^{\mathrm{nCh} \times \mathrm{nSamps}}$) belonging to the left or right-hand use,

$$\boldsymbol{f}(\boldsymbol{X}) = \left[\ \log(p(\,\mathrm{hand} = \mathrm{left}\,|\,\boldsymbol{X}),\ \ \log(p(\,\mathrm{hand} = \mathrm{right}\,|\,\boldsymbol{X}))\ \right]^T. \tag{4.8}$$

Since the probabilities of both hands are tied by $\sum_{\mathrm{hand}} p(\,\mathrm{hand}\,|\,\boldsymbol{X}) = 1$ only one of those probabilities are needed in the mapping, which makes the mapping equivalent to a $\mathbb{R}^1$ injection. However, we will continue with Eq. 4.8 as it enables the representation of the network as two separate hemi-networks. Both hemi-networks share the same parametrisation but only propagate activations from the corresponding half of the input (i.e. brain hemisphere) to the corresponding output (i.e. the ipsilateral hand).

Constraints are introduced in the architecture in the form of convolutional filter parameters that only allow interaction in the Hb, time and channel dimensions of the fNIRS signals. During training, zeroing-masks are applied to the fNIRS inputs affecting one randomly selected brain hemisphere with

equal probability. This hemispheric masking helps avoid solutions that rely exclusively on signals coming from only one hemisphere. In particular, the hemCNN implementation preserves the brain's hemispherical origin of information throughout all the layers and only relies on the difference between hemispherical brain activities to make a left or right-hand classification decision. A detailed description follows.

Design of the hemCNN architecture Figure 4.1 shows our CNN architecture design with the fNIRS input, 4 Convolutional Filter (CF) leading to 3 corresponding convolutional layers and one output layer with two output units representing both hands. The CFs convolve their respective input space leading to an output that is reduced in dimension depending on the filter size. By tailoring the shapes of the filters and their stride, we introduce constraints in the dimensions that the CNN can find relationships in. A CF is defined by its kernel shape (rows × columns) and its stride (rows × columns). Rows correspond to channels and columns to sampling time-steps. The kernel, $\boldsymbol{W} \subseteq \mathbb{R}^{m' \times n'}$, is the set of weights that the input is convolved with, and the stride is the number of steps the kernel shifts per row and column during the convolution. After the convolution an Exponential Linear Unit (ELU) function,

$$g_{ELU}(x) = \max(0, x) + \min(0, \exp(x) - 1) \tag{4.9}$$

is applied to each resulting unit. For an input $\boldsymbol{X} \subseteq \mathbb{R}^{m \times n}$, a CF defined by its kernel shape and stride leads to a preactivation,

$$\boldsymbol{Z} = \boldsymbol{W} \circledast \boldsymbol{X} + \boldsymbol{b} \tag{4.10}$$

which passed through g_{EL} results in the convolutional layer output,

$$\boldsymbol{Y} = g_{EL}(\boldsymbol{Z}) \tag{4.11}$$

with $\boldsymbol{b}$ expanding to the resulting size of $\boldsymbol{Z}$, the bias, and $\circledast$, the matrix convolution operator. Several of such functions, or convolutional layers, can be stacked, rendering a composite function.

We use discrete cross entropy $H(q,p) = \sum_{x \epsilon X} q(x) \log(p(x))$ between the real and predicted hand to optimise the weights of the network. As we advanced, the hemCNN architecture is equivalent to two parallel networks, or hemi-networks, mapping half of the input channels (i.e. brain hemisphere) to the probability of using the ipsilateral hand and sharing the same weights optimised for complementary

probabilities,

$$f_{\text{left}} : \mathbb{R}^{\text{nCh}_{\text{left}} \times \text{nSamps}} \longrightarrow \mathbb{R}^1 \tag{4.12}$$

$$f_{\text{right}} : \mathbb{R}^{\text{nCh}_{\text{right}} \times \text{nSamps}} \longrightarrow \mathbb{R}^1. \tag{4.13}$$

Since the network preserves the hemispherical origin of information in the outputs, it learns hemispherical brain features that provide higher output values for the ipsilateral hand than for the contralateral one. For example, a right hand example is correctly classified if $f_{\text{right}} > f_{\text{left}}$, and vice versa for the left hand.

In Table 4.1, the particular hemCNN kernel sizes and strides are shown. Namely, for an input arranged as in Fig. 4.1, CF1 convolves two input rows (HbO and HbR of the same channel) into one row dimension of the resulting CL and 10 input time-samples into one time-convolved output. Once the column convolution of the first two rows (i.e. $\text{HbO}_{\text{Ch1}}/\text{HbR}_{\text{Ch1}}$) is finished, the row stride of 2 makes the kernel shift to the next $\text{HbO}_{\text{Ch2}}/\text{HbR}_{\text{Ch2}}$ pair instead of to the $\text{HbR}_{\text{Ch1}}/\text{HbO}_{\text{Ch2}}$ pair of two consecutive but different channels. The latter would significantly undermine learning as we would be forcing the CF to filter signals of different nature, and the architecture only depends on this single filter. Since no padding is used in the convolution, the time and channel axes are compressed convolution after convolution.

Training, validation and testing The fNIRS time-series are z-scored per trial using the mean and standard deviation of the whole channel set per Haemoglobin (Hb) type. All channels are used for classification. In contrast to standard approaches, a baseline correction is not performed, preserving time and inter-channel relationships while enabling the conditioning of the input to the CNN as it is standard practice. Training examples are obtained from fNIRS trials and augmented by (1) randomly cropping 19 s time windows from an initial $0 - 35$ s trial window starting at the go-cue and (2) by multiplying each trial by a random number drawn from a uniform distribution between 0.5 and 1.0.

hemCNN	CF1	CF2	CF3	CF4
kernel $(m \times n)$	2×10	1×5	1×5	12×5
stride $(m \times n)$	2×1	1×2	1×2	12×5
Total	94			

Table 4.1: Filters shapes and the total number of weights for hemCNN. Kernel shape (rows × columns) and stride (rows × columns), with rows corresponding to interleaved HbO/HbR channels and columns to sampling time steps. Only one filter is applied per layer. Kernel shapes and strides preserve the hemispherical localisation. The total includes bias weights, one per each CF.

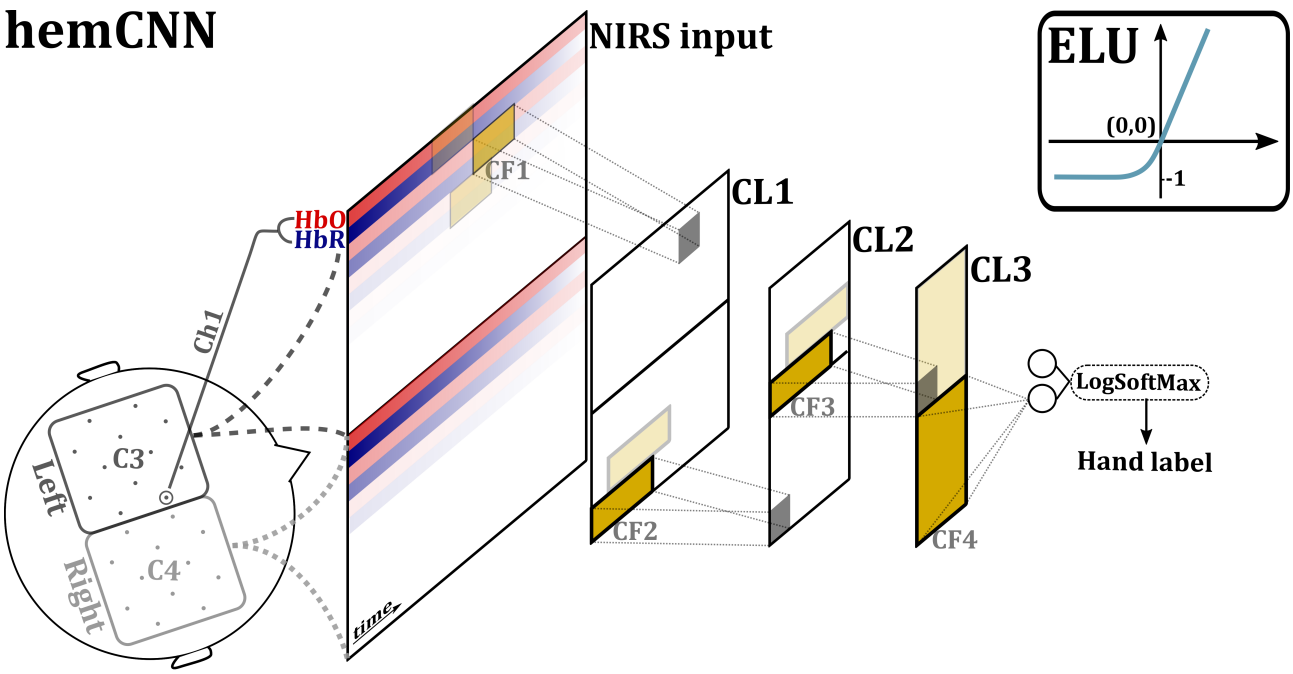


Figure 4.1: HemCNN architecture. From left to right: Channels in the input are arranged per row as HbO/HbR consecutive pairs with the upper half of the rows corresponding to the left hemisphere and the one at the bottom to the right hemisphere. Columns correspond to time steps. Each convolutional filter (CF) renders a deeper convolutional layer activation (CL) after convolution. The last layer consists of a unit per hemisphere whose activations pass through a LogSoftMax function. The unit rendering the highest value after the application of this function is selected as the label of the hand: left for the upper unit, and right for the one in the bottom. The exponential-linear function is used as non-linearity for all units except the output units.

In the first case, we want to encourage hemCNN to locate the features at any time during the trial (Schirrmeister et al., 2017). In the second, we further discourage the network from depending on absolute amplitude values for classification. The dataset is augmented 5 times by randomly applying these transformations to the original set.

A challenge with binary classification in our context is that by the task's design, and thus the data, the classification into mandatory "left" or "right" options may result in a confounding outcome. Namely, a network that may learn "left" vs "non-left" instead of "left" vs "right". While from a decoding correctness perspective, this is fine and often overlooked in BCI classification, the goal of this work is to train an explainable deep learning architecture. Thus we want a network that can reliably capture the distinctive activities evoked by left or right-hand use across hemispheres. One option to overcome this would be to include a third "reject" option, i.e. not left and not right hand (Trakoolwilaiwan et al., 2017). However, the generation of a suitably diverse set of brain activities to drive this reject option non-trivially (i.e. not with resting-state brain activity) is complicated in itself, and there is no guarantee that this "reject" option would be learnt as such. In other words, the actual reject option could end up being the "right" class, with everything not fitting into the "left" and "reject" classes classified

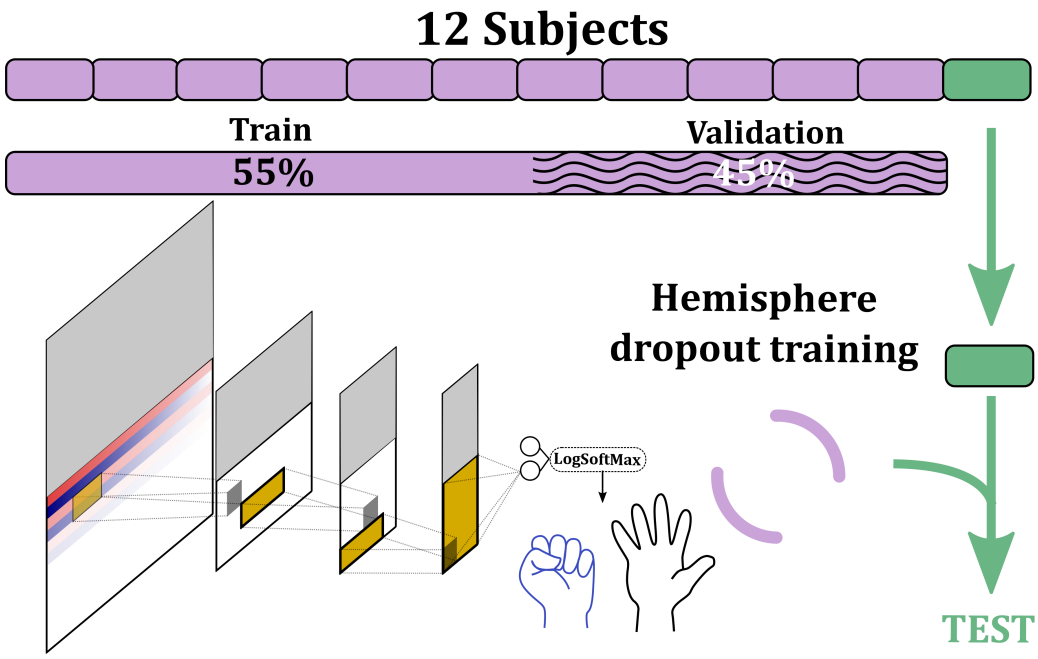


Figure 4.2: HemCNN leave-one-subject-out training. The data from twelve ($N = 12$) subjects are used to train models, leaving one subject out for testing. In the architecture (bottom), grey shaded areas depict the masking of a hemisphere for each example. The left or the right brain hemisphere was masked or dropped-out with equal probability for left and right-hand examples. In the figure, the left brain hemisphere is dropped out during the training of a left-hand example.

as “right". We here use regularisation. Specifically we drop-out (Srivastava et al., 2014) input units reflecting either hemisphere with equal probability (Fig. 4.2). Namely, for each fNIRS example $\boldsymbol{X}$ fed into the network, either the left or the right hemisphere are randomly zeroed, with a balanced amount of left and right hemispheres zeroed in each batch. When we apply this second approach, we consistently obtain network weights that better reflect the contrast between hemispherical activities for each hand that otherwise only appeared in a few trained models. We use the cross-entropy loss as the error function and the Adam optimiser (Kingma and Ba, 2015) with 10^{-3} weight decay. The training batches contain 10 left/right balanced shuffled examples, and the learning rate starts at 0.03. Training is iterated for 15 epochs (i.e. 15 times over the augmented training-set), with the learning rate decaying at a 0.9 rate per epoch.

Leave-one-subject-out cross-validation and testing Leave-one-subject-out testing is a particularly demanding generalisation condition typically avoided in BCI work, as it requires the decoding algorithm to operate on the unseen subjects (Xiloyannis et al., 2015), not just unseen trials as commonly used. Leave-one-subject-out testing measures the robustness of the model’s performance against changes between subject physiology, brain anatomy, sensor placements, experimental conditions, subject behaviour and all remaining variables specific to each subject in particular (Fig. 4.2).

For each subject, the 20 first balanced left/right-hand trials are used as the training set, and the remaining, if applicable, are discarded. For each left-out-subject, 55% of data from the remaining subjects is used for training and 45% for validation. Training and validation data are split randomly every run, and 10 runs per left-out subject are executed. Training data is augmented in the way mentioned above by running 5 times per training example the combined augmentation techniques (random time crop and random multiplicative noise). The training procedure is run 10 times, delivering 10 different models per left-out subject. Each model has a different random initialisation and trains for 15 epochs. In the end, every run is trained with $11 \times 20 \times 0.55 \times 5 = 605$ augmented examples. The 99 validation examples are used to select one model out of the 10, and the remaining 20 belonging to the left-out-subject to produce all results per subject using the selected model.

HemCNN model selection and analysis Out of the 10 trained models per subject, the model with the highest validation accuracy is selected to analyse the method's accuracy and learned features. Both the accuracy and learned features are analysed using the unseen subject test data, and the test results are aggregated across subjects to provide statistical measures. We first analyse the Hb relationship (i.e. the relationship between HbO and HbR) by studying the first convolutional filter (CF1) and then use relevance propagation (Bach et al., 2015) to study the patterns that across models and subjects are more stable and relevant for the classification.

Haemoglobin filter interpretation We are interested in understanding what Hb relationships elicit higher activations in the network. To do this, we use pairs of centred sinusoids representing HbO and HbR and filter them with the corresponding CF1 layer of each trained model. The bias of CF1 is removed since the sinusoids are centred around zero.

To check if the response of the filter is sensitive to delays between HbO and HbR we also study the filter response at different time lags.

$$\begin{aligned} y_{\text{HbO}} &= A\sin(2\pi f t) + \epsilon \\ y_{\text{HbR}} &= A\sin(2\pi f (t+\tau)) + \xi \end{aligned} \tag{4.14}$$

The frequency, f, is set to 0.1 Hz as representative of the HR. Amplitudes are the absolute value of the samples drawn from a standard normal distribution, $A \sim \mathcal{N}(0,\ 1)$. Independent noise is added to each sinusoid $\epsilon, \xi \sim \mathcal{N}(0,\ 0.5)$. The CF1 output is resampled to the input sampling rate by interpolating

the output time-compressed points between the original time points. This interpolation allows us to align the input and the output and observe what input relationships lead to higher CF1 responses. Delays range from $\tau = -5\,\text{s}$ (in phase) to $\tau = 0\,\text{s}$ (in antiphase). The procedure is repeated 1000 times, and averages are computed.

Relevance propagation for model interpretation To further investigate the nature of the features discovered by hemCNN in the fNIRS brain signals, we use relevance propagation (Bach et al., 2015). This technique allows us to inspect what units in each layer are relevant to produce a particular output for a given example. The standard procedure starts by feeding an example to the network and zeroing all output units except the one representing the class for which we want to find relevances. This output's value is backpropagated so that each of the activations in previous layers is given a relevance weight in the generation of the observed output (Bach et al., 2015). We can then observe the relevance of each unit when choosing the category represented by the selected output. Note that all non selected outputs only backpropagate zeros and do not affect the selected category's relevance.

We adapted this procedure to hemCNN. Recall, hemCNN can be seen as two hemi-networks (Eq. 4.12), each observing a brain hemisphere and producing an output representing a hand decision. For correct left-hand decoding, the left-hand output should be higher than the right-hand and vice-versa. If we were to follow the standard procedure, we would need to mask the left-hand output by zeroing it. However, this would backpropagate zeros through all the left hemi-network layers defeating our purpose of revealing both ipsi and contralateral patterns associated with each hand. To solve this problem we backpropagate the value $|\,p(x) - \frac{1}{2}\,|$ for both outputs. We here subtract $\frac{1}{2}$ to reference the output to the indecision level. By doing so, we make the indecision level the only case where zeros are backpropagated to both hemispheres. The absolute value is used to get only positive values as required for relevance propagation. Since correct left-hand decoding requires higher left-hand output values than right-hand ones, we only consider the backpropagation of those layer activation values relevant to producing lower contralateral outputs and higher ipsilateral outputs, i.e. those driving correct classifications. This amounts to use the negative and positive relevance propagation variation of the procedure in Bach et al. (2015) with parameters set to $\beta = 1$ and $\alpha = 0$ for the contralateral output and $\beta = 0$ and $\alpha = 1$ for the ipsilateral.

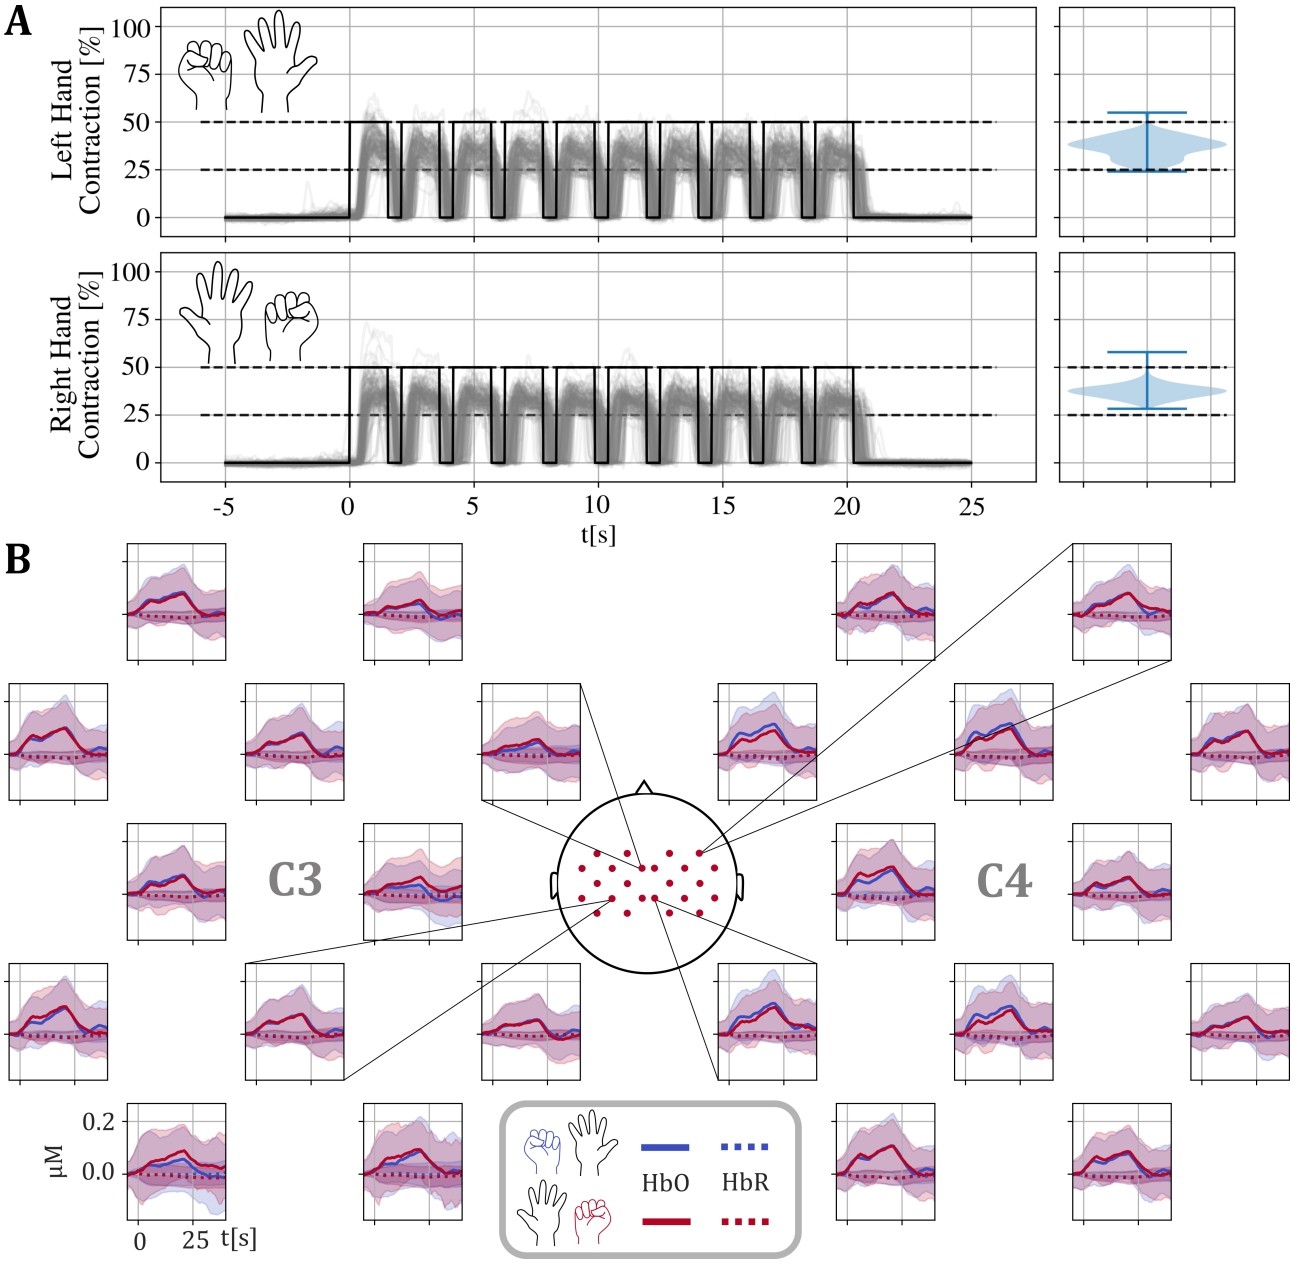


Figure 4.3: **A** Voluntary contraction per hand condition for 12 subjects. On the left, the distribution of the 90^{th} percentiles values per trial. Contractions are inside the expected range (25 – 50% MVC). **B** Evoked Haemodynamic Response (HR) average per hemoglobin specie (HbO and HbR) and hand condition. A global response can be observed with delayed start and decay. The high variability is depicted by the high standard deviation of the signals across trials. Compare with simulation in Fig. 2.6.

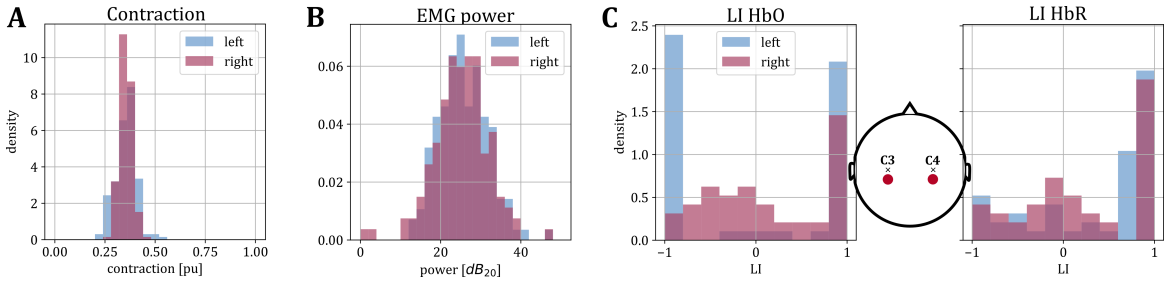


Figure 4.4: **A** Contraction levels (90^{th} percentile of each trial) for the left and right hand with no significant differences (Mann-Whitney rank test, p = 0.15). **B** EMG power levels (90^{th} percentile of each trial) for left and right-hand with no significant differences (Mann-Whitney U test, p = 0.11). **C** LI between posterior to C3 and C4 fNIRS channels for HbO (left) and HbR (right). Although some differences can be observed in the LI distribution across hands, the variability of the signal leads to no significant statistical differences for HbR (Mann-Whitney U test, $p = 0.41$) and HbO (Mann-Whitney U test, $p = 0.07$).

4.3 Results

To evaluate the performance of hemCNN resolving the HR in fNIRS we used the data of 12 participants from HYGRiP performing a left/right hand-grip task. During the task, the force, the arm EMG and the brain activity in the form of the cortical Haemodynamic Response (HR) from fNIRS were continuously recorded (Fig. 4.3A and B). Breathing was also recorded to account for artefact sources. Figure 4.3B depicts the high variability that the task evokes in the HRs which any classification algorithm needs to confront. HemCNN is trained to detect what hand (left or right) is being used from these HRs as in a BCI decoding paradigm.

We first validate that the selected subjects correctly performed the motor task using contraction and EMG activities during the task (Fig. 4.4A and B respectively). Both contraction and EMG activity data show that all participants performed the left/right hand-grip task correctly. Figure 4.4A and B show that both hands engaged similarly in both measures: The contractions levels (90^{th} of each trial) are in the 25 – 50% MVC range (Fig. 4.4A). The right and left-hand contraction levels show insignificant differences (Mann-Whitney U test, $p = 0.15$). Similarly insignificant differences are found for left and right EMG power (Mann-Whitney U test, $p = 0.11$, Fig. 4.4B). These results show that the variability in the inter-hand, the inter-trial and the inter-subject force and muscle activities are not likely to play a role in potential differences in cortical HR activity. In contrast, differences in cortical activities might be specific to the hand (left or right) that the subjects used.

4.3.1 Standard analysis

Consequently, we then ask if cortical HR activity shows hemispherical differences associated with each hand. To do so, we use the LI shown in Fig. 4.4C. The LI aims to enhance hemispherical differences and determine which hemisphere is more active during a particular condition. We want to see if LIs exhibit different hemispherical dominance associated with the left and the right hand during force production. When LIs are computed between channels 7 and 19 the differences between the left and right hand hemispherical dominance are not significant (Mann-Whitney U test, HbO p = 0.07 and HbR p = 0.41). Although visual differences can be observed in Fig. 4.4C for the LI distribution, the high variability in the signal may reduce the significance of the differences. In fact, these levels of significance are susceptible to the selection of LI computation parameters, such as the baseline correction times. Although this approach fails to find consistent cortical activation differences, the sensitivity to LI parameter selection and the visual differences in the empirical LI distribution suggest that differences may exist but are not adequately captured by these traditional ΔHb features.

4.3.2 HYGRiP benchmark

If hand-specific differences in cortical haemodynamic activity exist, the previous standard methods might be unable to find them and more sophisticated feature extraction methods like GLMs, General Linear Model with Hierarchical Principal Component Analysis (GLM-hPCA) and, mainly, CNNs might provide better results. The ability to distinguish between HRs evoked by the left or the right-hand grips can be measured as the hand classification accuracy of these methods using HYGRiP's fNIRS signals.

We have described the technical arrangements of the task in detail and made the dataset publicly avail-

Decoder	Class. Acc. [%]
hemCNN	**77.7 (65.4 - 84.7)**
hemCNNo	**77.9 (63.5 - 82.6)**
hemCNNr	**69.3 (62.5 - 84.2)**
GLM	55.0 (51.9 - 60.0)
GLM-hPCA	54.4 (50.0 - 63.3)
ΔHb	54.4 (46.2 - 60.0)
LI	50.0 (45.6 - 52.5)

Table 4.2: HYGRiP benchmark performance for fNIRS-based hand classification accuracy during the production of strenuous force, median (25^{th} - 75^{th} percentile). In bold, significantly outperforming methods compared to non-hemCNN approaches (Kruskal-Wallis multiple comparison, Tukey corrected, α = 0.05). HemCNNo and hemCNNr stand for hemCNN using only oxygenated or deoxygenated haemoglobin, respectively.

able and ready to exploit in the HYGRiP dataset (Ortega et al., 2020). Finally, we made hemCNN's code available so that others can benchmark any novel decoders with ours (available here). In this sense, we believe that the task, the collected data, and hemCNN's performance set a new benchmark for fNIRS-derived BCI decoding - the HYGRiP benchmark.

The specific benchmark values are the median (25^{th}-75^{th} percentiles) of the leave-one-subject-out test accuracy for hemCNN, GLM, and GLM-hPCA along with that of the traditional ΔHb and LI features (Table 4.2 and Figure 4.5). Except for hemCNN, the remaining methods use a simple-tree that classifies the respectively computed GLM, GLM-hPCA, ΔHb or LI features into the used hand.

As we can see, hemCNN methods outperform GLM, GLM-hPCA, ΔHb and LI approaches (Kruskal-Wallis test, Tukey correction for pairwise comparison, $\alpha = 0.05$; Figure 4.5). The use of both HbO and HbR compared to only using one of them does not lead to significant differences, although the absence of HbO in hemCNNr leads to a drop in accuracy (median accuracies, hemCNN 77.7%, hemCNNo 77.9% and hemCNNr 69.3%, $p > 0.9$).

We also used a standard CNN with 4 convolutional filters per convolutional layer and 3 fully connected layers to classify the signal (not shown). This standard CNN achieved a median accuracy of 80.3 (69.0 – 86.1)% which was not significantly different to that of hemCNN, 77.7 (65.4 – 84.7)% (Kruskal-Wallis, Tukey corrected test, $\alpha = 0.05$).

Non hemCNN methods, except LI, show similar performances around 54.4 – 55% median accuracies. In particular, the addition of hierarchical PCA in GLM leads to the highest performance in non hemCNN methods for subject 1 (accuracy 75%). LI performs with a random median accuracy level (50%). This is in line with the low significances of the first approach, which might have been further hindered by the increase of dimensionality after including all channels. Interestingly, non hemCNN approaches show at least one case of classification accuracies below the random level (< 50%). This indicates that some subjects' classification rules do not generalise to another one and may even act in an opposite classification direction.

4.3.3 HemCNN's interpretation

Finally, we show how, in addition to providing higher accuracies, the hemCNN learns HR features with physiological interpretation from fNIRS. We analyse hemCNN with HbO and HbR since it is more interesting from a physiological perspective. In particular, the analysis of the activations across

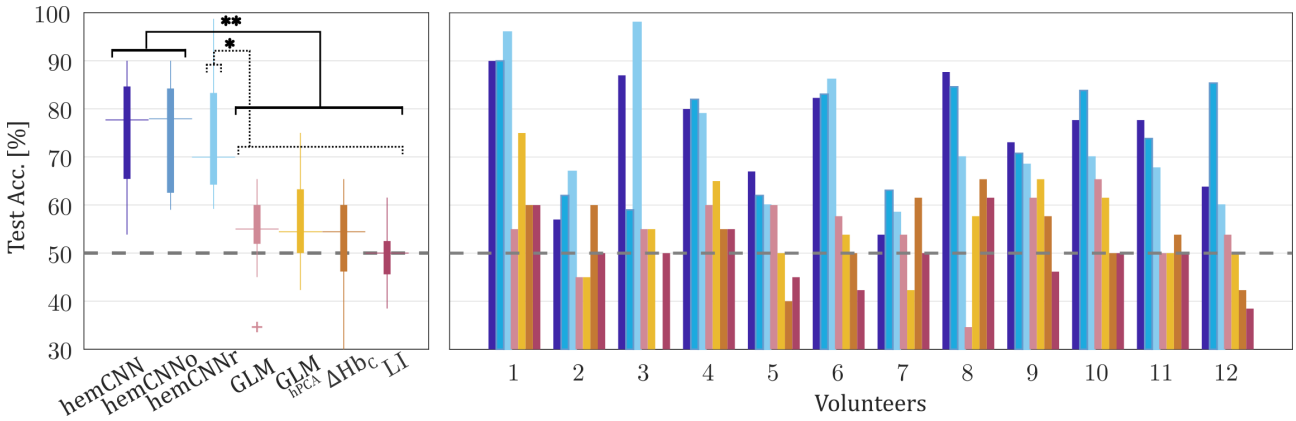


Figure 4.5: Leave-one-subject-out test accuracy median across subjects (left) and per subject (right) for the best validation model out of 10 trained models. Each model is trained in all subjects but the left-out one, in which the model is tested. The ordering of methods shown on the left is preserved per subject on the right. Cold colours denote CNN architectures: hemCNN with HbO (hemCNNo), HbR (hemCNNr) or otherwise both inputs (hemCNN). Warm colours denote standard methods to extract features from HbO and HbR which are classified using a simple tree: general linear model (GLM) features, GLM with hierarchical PCA features (GLM-hPCA), ΔHb features and LI features. Left, hemCNN performs significantly better ($p < 0.01^{**}$ for hemCNN and hemCNNo and $p < 0.05^{*}$ for hemCNNr) than standard approaches. Right, also on a subject by subject basis hemCNN methods outperform all standard decoding approaches.

hemCNN layers reveals different spatio-temporal structures in the cortical activation for each hand which we dissect hereafter.

To understand what the learned Hb relationships consist of, we look into the first convolutional filter (CF1) of hemCNN (Fig. 4.6A). This filter convolves HbO/HbR pairs (from the same channel) in time. We feed sinusoids (Eq. 4.14) to this filter to analyse its response under certain HbO/HbR delay conditions. Note that training hemCNN with a leave-one-subject-out cross-validation approach produces 12 different hemCNN parameter-sets for the neural network weights. The 12 parameter-sets produce a consistent overall interpretation which is best explained by looking at the individual cases. Five of these parameter-sets, present the highest positive activation at HbO/HbR cross-over points where HbO increases and HbR decreases as shown in Fig. 4.6A. The high activation for this feature is stable for HbR delays in the $(0, 2])$ s range. Another three parameters sets showed the same but the peaks corresponded to the highest negative activation (the activation was inverted). Two have the highest activation in the HbO/HbR peak separation. The remaining two show a response where the positive activation is highest between the cross-over point and the separation peak. Of all solutions, pure HbO/HbR cross-over are learned in 61.5% of the models, suggesting that such a feature is the most stably present in the data. The models maintain a high level response for different HbR lags in the $(0, 2])$ s range. Other solutions are also closely related to this observation. The existence of this stability across this range of HbR delays suggests that the filters are stable against variability in the signal that artefacts and non-stationarities could introduce.

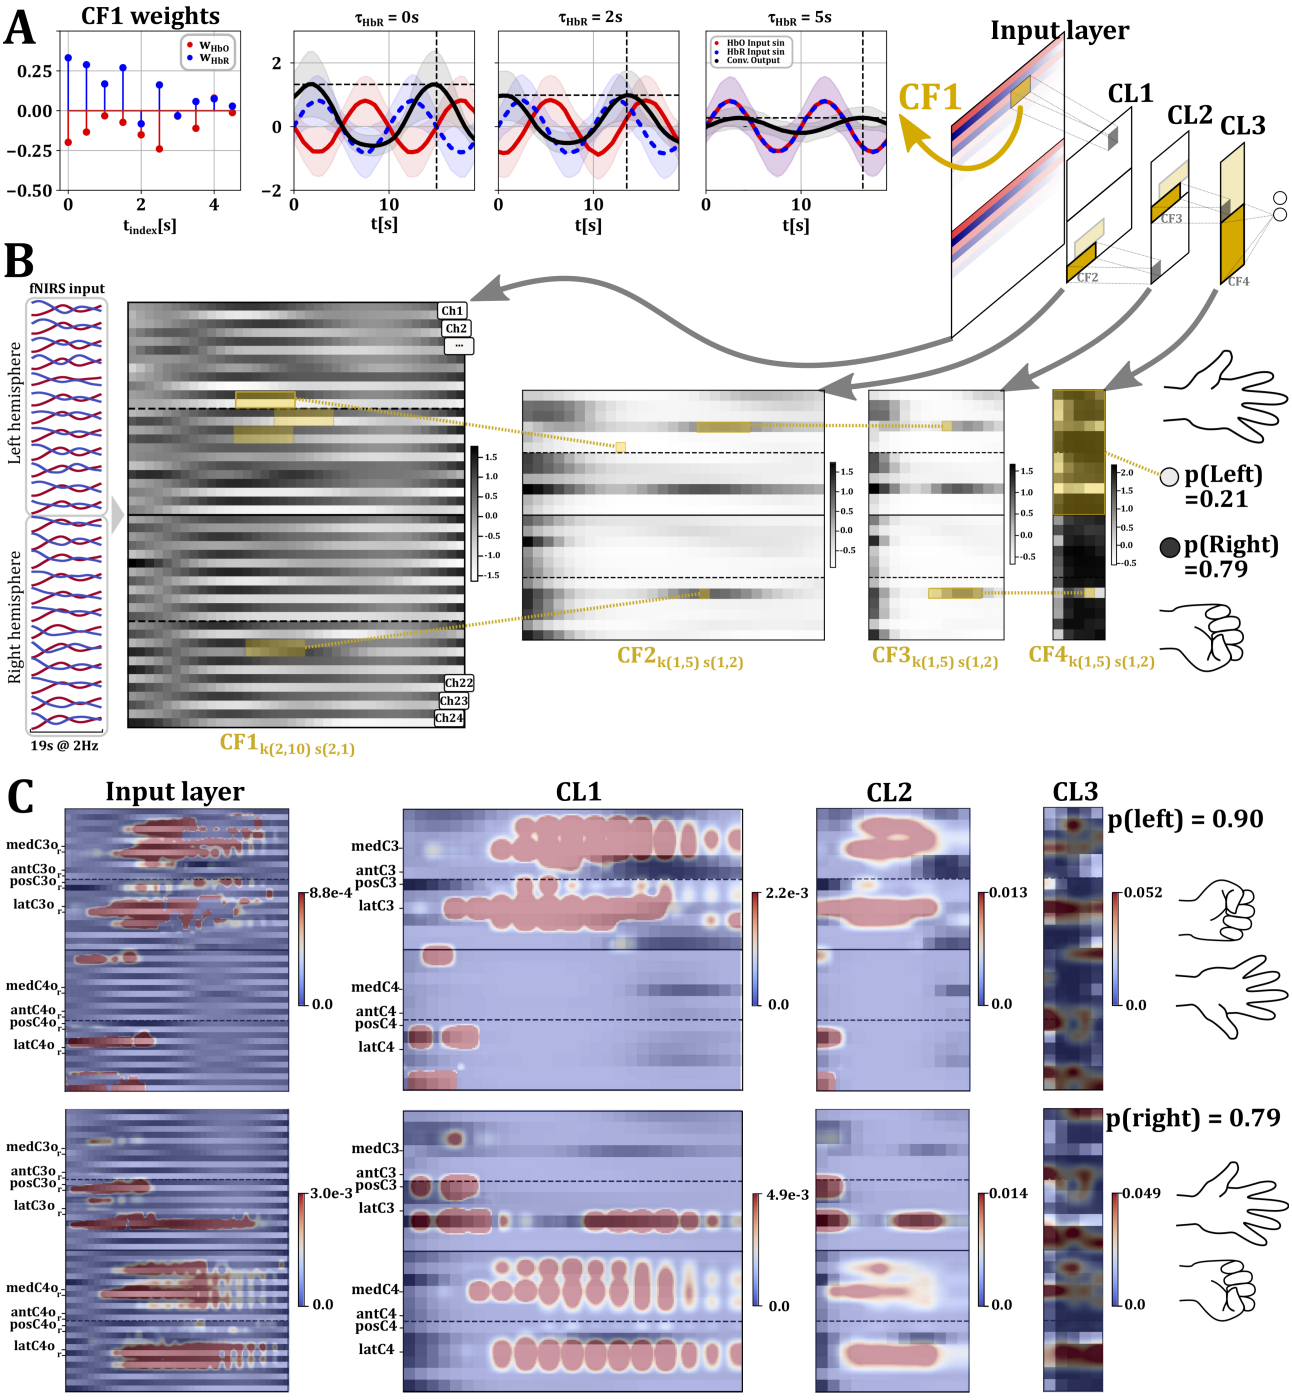


Figure 4.6: HemCNN representative interpretation. **A** First filter (CF1) response. Left, representative CF1 weights associated with HbO and HbR. Right, CF1 response for different HbR delays. Delays in the $0-2$ s range render high filter responses close to HbO/HbR cross-overs. **B** Unit activations throughout hemCNN for one fNIRS example. Left to right, fNIRS input and convolutional layers (CL) for one right-hand trial. In yellow, the four convolutional filters (CF) and their size. The filters preserve the hemispherical origin of the information. Inter-hemispheric differences start being evident in CL1 and more apparent in deeper layers. **C** Relevance per layer for two representative examples. Most high probability examples showed higher relevances in anterior and posterior to central contralateral channels and medial and lateral to central ipsilateral channels.

If we now move deeper into the hemCNN architecture, we can observe that the convolution of fNIRS inputs with CF1 results in a time-compressed output with a single row per fNIRS channel. Deeper layers further convolve these outputs in the time dimension (CF2 and CF3) and, finally, in the channel and time dimensions (CF4). Each layer revealed interesting spatio-temporal structures for each hand that we describe next. Figure 4.6B shows a representative case of such spatio-temporal structure in one fNIRS example. The structure can be directly observed by looking into the convolutional layer activations. The first column on the left represents the HbO/HbR arranged in pairs per channel (fNIRS input, z-scored to the mean and standard deviation of values in the same sample per Hb type). The next column represents the same signals but visualised as an array with amplitudes coded into a grey scale. The following columns present the convolutional layer activations in the same way. The convolutional time-compression effect can be observed as the narrowing of the convolutional layers' activations in the time axis from left to right. For CF1, we can see how areas with high HbO/HbR contrast lead to negative responses and areas with less contrast and close to cross-overs lead to positive responses. CF4 convolves each hemisphere to a single unit that, after normalisation (in this case, SoftMax for readability), leads to a probability of the example belonging to the left or the right hand. An asymmetric pattern can be observed in CL1 and deeper layers. In most examples, similar asymmetries can be observed.

However, there are two sources of ambiguity in this interpretation. First, the observed activations might be present but not relevant to the final decision. Moreover, second, across left-out subject models, the solutions can be "symmetrical". Namely, we can multiply by -1 each CL and CF and still have the same classification decision. In this case, we lose the sense of what a negative or positive hidden unit activation entails in the decision when comparing models that have reached "symmetrical" solutions. We overcome these by analysing the Relevance Layer Propagation (RLP). Our RLP implementation highlights the unit activations that drive the classification decision and are always positive across layers regardless of the model. Figure 4.6C shows the relevant activation areas across layers for two examples corresponding to different hand use. Asymmetries across hemispheres and hand use can be observed. Interestingly, relevance in the contralateral hemisphere appears earlier than in the ipsilateral with a preference for channels that are anterior or posterior to the central channel (APc) rather than channels that are medial or lateral to the central channel (MLc)[5]. The average across such relevances

[5]We use here anatomical nomenclature. For a given point in the body, "anterior" or "posterior" describes the direction towards the front or back of the body, respectively. "Medial" describes the direction towards the sagittal plane, i.e. the plane dividing the body into two symmetric, right and left, halves. "Lateral" describes the direction away from the sagittal plane.

are consistent with these patterns, but variability makes time differences vanish, and only channel asymmetries can be observed.

To achieve a population-level analysis, we computed the absolute value of the correlation between the relevance and the activation of each unit across examples. This analysis revealed more evident differences in the spatio-temporal cortical activation structures associated with each hand. Figure 4.7 shows the results ($p_{\text{correlation}} < 0.01$, units with less significance than this are zeroed). The first two columns on the left show the correlation for each hand. The third column shows the difference between the previous columns. Each row represents a different hemCNN activation layer. In the input layer, the lack of definition of an HbO/HbR feature makes correlations sparse. Nonetheless, differences across hemispheres can be appreciated. In the next activation layer, CL1, CF1 now defines the HbO/HbR feature when filtering the input and enables greater consistency between relevance and activation. This feature makes antisymmetric patterns already distinguishable, with early ipsilateral relevance in APc channels compared to later contralateral relevance in MLc ones (row 2, column 3 of Fig. 4.7A). The time compression effect of the following filters (highlighted as an equivalent sampling frequency under activation layer titles on the left) allows for an alignment of the previously found features and much clearer differences across hemispheres and hand conditions. In particular, the topographic projection of CL2 correlations (bottom of Fig. 4.7B channels around C3 and C4 highlighted in purple) shows a response that starts with a strong early response with already observable hand specific hemispheric patterns ($0-3\,$s). These patterns become clearer later ($9-15\,$s) and end vanishing ($15\,$s onward).

Finally, to discard that cardiorespiratory variations drive these features, we take two measures. The first consisted of band-pass filtering the fNIRS signals below $0.3\,$Hz, well below the heart-rate $\sim 1\,$Hz, ensuring that the pulse was filtered out. The second consists of training an architecture equivalent to hemCNN, with the same number of layers and filter shapes, but adapted to have the breathing signal as its input (one single channel, low-pass filtered below $1\,$Hz and downsampled to $2\,$Hz). The resulting models' median accuracy (50.86%) was very low and not significantly different from the 50% random classification accuracy (sign test, $p = 0.078$), discarding the possibility of breathing rhythms potentially informing on what hand was used. The two measures together rule out that these systemic artefacts inform our hand decoding task and learned features. This confirms that we are indeed able to decode and track cortical haemoglobin changes evoked by the left and right-hand grips.

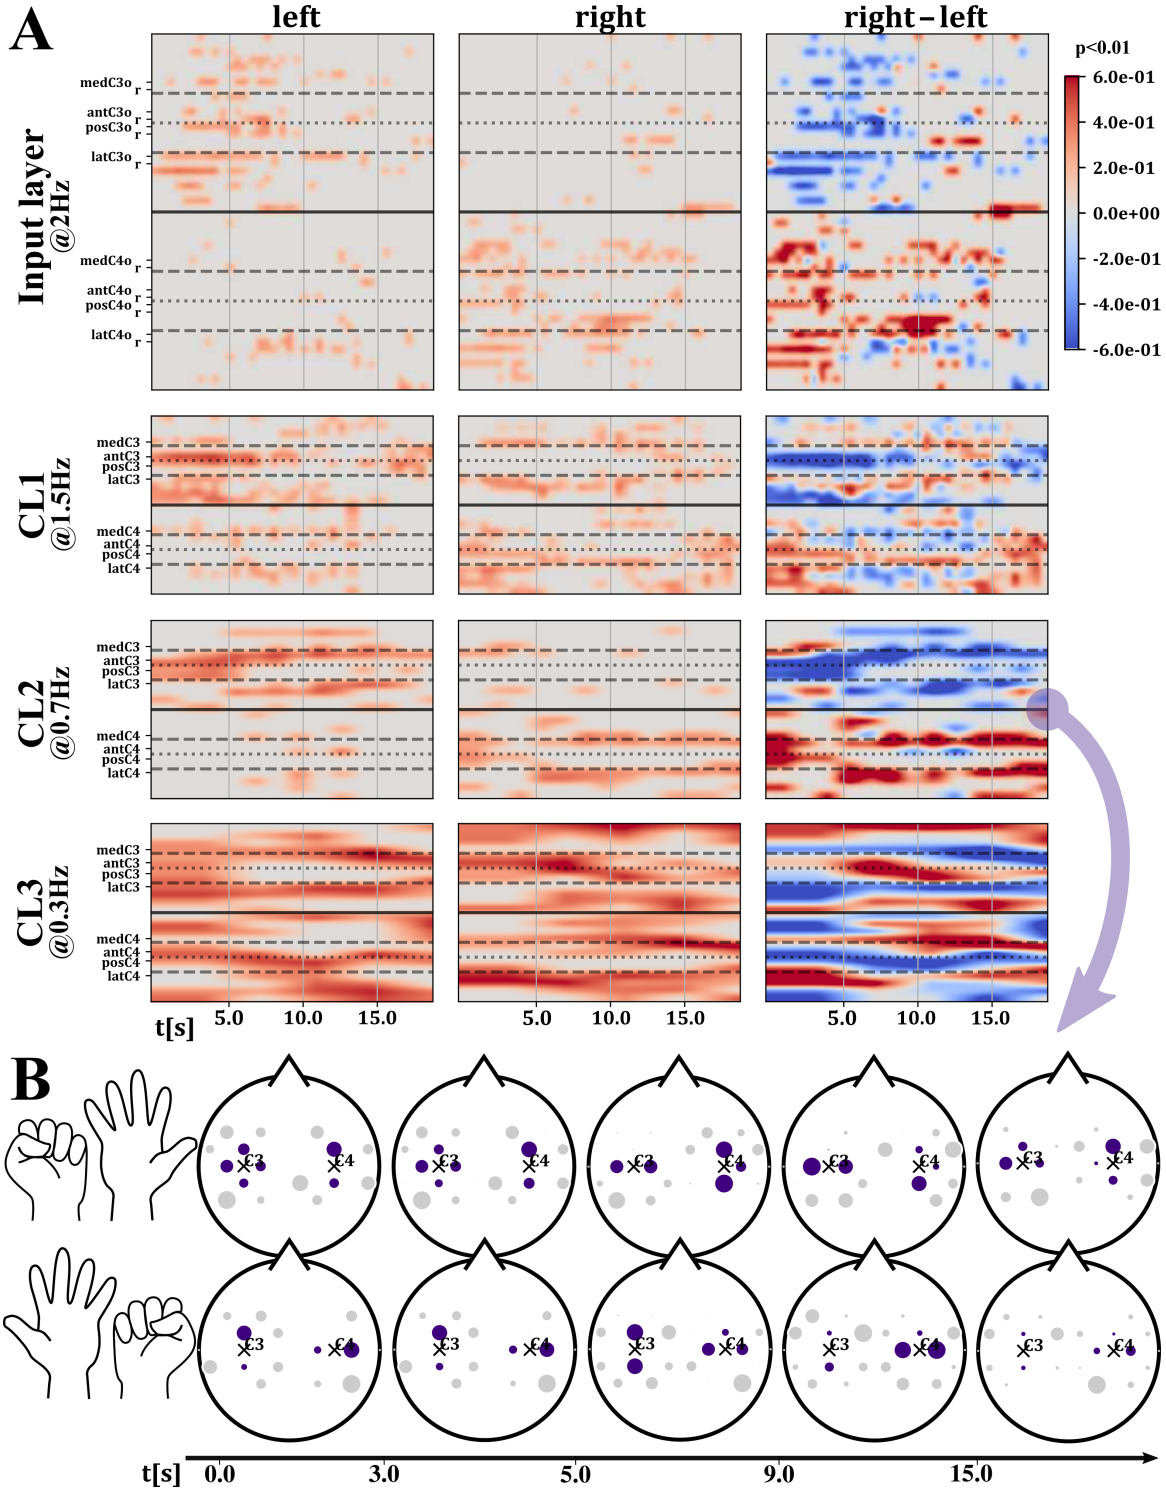


Figure 4.7: HemCNN cross-subject interpretation. **A** Absolute value of the correlation between units activities and relevances. Correlations are presented for the left and the right hand and their difference (columns); and for each hemCNN layer (rows). Layer activations are upsampled and aligned in time to the input layer, with time resolution decreasing with depth. The lateralisation of activities is better observed the deeper in the architecture due to the application of convolutional filters. **B** Topographic projection of the correlations in the activation layer CL2. Clear hemispherical differences associated with each hand appear during the task that were hard to identify in the raw signal averages (Fig. 4.3B)

4.4 Discussion

We used the motor cortical Haemodynamic Responses (HRs) evoked by a left/right-hand grip task from HYGRiP. We also used several signals from our HYGRiP dataset that reflect the task behaviour, like EMG, or confounding factors such as breathing. In this chapter, we presented and used a novel Deep Learning (DL) approach, hemCNN, to perform decoding and neural data analysis. We also made the results interpretable through a twist in the design of the hemCNN architecture. As intended, the left/right-hand gripping task-evoked highly variable and non-stationary hemodynamic responses. Using these responses, hemCNN outperformed popular methods, especially GLM, at decoding the hand used. Using hemCNN, we also revealed spatio-temporal hemispherical differences during the whole duration of the task that were not resolved in (Derosière et al., 2014) or took more than 30 s to appear in (Shibuya et al., 2008, 2016). Indeed, the better performance of hemCNN is based on the better detection of complex and variable Haemodynamic Response (HR) features learned from the data. By designing the DL architecture, we could open the black-box and interpret these features scientifically.

Strenuous hand grips evoked a highly variable HR To test the limits of our decoding approach, we devised a challenging task in HYGRiP (Ch. 3) that required a high contraction range output (25 – 50% MVC), fast pace and visual feedback. These factors previously challenged the decoding of fNIRS signals (Shibuya et al., 2008; Derosière et al., 2014; Wriessnegger et al., 2017). We also introduced fast switching on/off dynamics (1 Hz) in the grasping execution and shorter resting times and stimulations. As intended, the evoked Haemodynamic Responses (HRs) had high variability and were apparently symmetrical over the sensorimotor cortex. Namely, no HR differences were associated with the left or the right hand using conventional approaches (Fig. 4.3 and Fig. 4.4C). Additionally, we verified that both hands were similarly engaged in terms of EMG and force production during the task (Fig. 4.4A and B). In this way, we rejected that potential differences in cortical activations were due to how the hands engaged in the task.

HemCNN improves the classification of highly variable HRs We chose not to make any assumptions on the HR features and developed a CNN for fNIRS time-series decoding called hemCNN. HemCNN performs significantly better than ($\sim 78\%$, $p < 0.01$, Fig. 4.5) the General Linear Models (GLMs) features (Ye et al., 2009) and a more powerful variant of GLMs I implemented with PCA

(see Fig. 4.5). Other commonly used features, such as amplitude changes (ΔHb) and Lateralisation Index (LI) led to decoding at almost chance level (see Fig. 4.5). When a standard CNN was used for classification, we achieved a slightly higher median accuracy (80.3%) which was not significantly different to that of hemCNN, 77.7% (Kruskal-Wallis, Tukey corrected test, $p > 0.05$). However, in contrast to hemCNN, the standard CNN architecture does not guarantee that the classification is due to interhemispheric differences. Moreover, the standard CNN does not allow to trace back each hand output activity unequivocally. Instead, in the standard CNN, these activities are sparsely distributed across the different filters and fully connected units. These two aspects defeat the application of standard CNN to prove our neuroscientific hypotheses. Namely, that (1) interhemispheric differences in HR are related to hand-specific cortical representation measured with fNIRS, and (2) that the decoder can be fully interpretable. We did not include the standard CNN in the principal analysis for these reasons. Nonetheless, improvements in the decoding accuracy can be expected if architectures with more parameters are used. In summary, our analysis demonstrates that interhemispheric differences in fNIRS signals exist across hands but require appropriate feature learning to be detected.

HemCNN provides a physiological interpretation We designed hemCNN to be interpretable, i.e. it allows us to open the black-box models used and physiologically interpret the learned features. We found that the hemCNN mostly learns to detect the simultaneous increase of HbO and decrease of HbR (Fig. 4.6A), which hereafter we refer to as "HbO/HbR functional signature".

Moreover, the low CF1 filter response for HbR delays of more than 2 s (Fig. 4.6A) suggests that hemCNN filters out systemic artefacts. This response is in line with previous work associating systemic artefacts with that range of HbR delays (Boden et al., 2007), and justifies the use of both measures (beyond mere physiological interest) to improve the interpretability of functional activations regardless of systemic variations, without the need to invest in additional short-distance channels (Zhang et al., 2009; Gagnon et al., 2011).

The use of only one HbR in the input led to a drop of the decoding accuracy (Fig. 4.5) that is in line with previous studies showing that HbR changes are generally less consistent across subjects than HbO (Maki et al., 1996; Miyai et al., 2001).

Further exploration of the learned HbO/HbR features (detected by the first convolutional filter) in the input layer reveals that the HRs are sparse in time across trials (Fig. 4.7A, row 2). The time compression effect of the two subsequent convolutional filters enables the temporal structure to be

preserved across trials and results in higher correlations at the population level (Fig. 4.7A, rows 3 and 4). While most studies resolve the hemodynamic response in the $0.01-0.3\,\mathrm{Hz}$ band we can increase the resolution to $\sim 0.7\,\mathrm{Hz}$ (Fig. 4.7A, row 3), which is similar to the estimated auto-regulation constant of the HR balloon model ($0.65\,\mathrm{s}^{-1}$ (Buxton et al., 1998; Friston et al., 2003)).

The detection of the HbO/HbR functional signature and the alignment in time across examples enabled the last layer to learn channel relationships associated with each hand's use. Finally, this allowed to reveal a consistent hand-specific hemispherical activation early at the beginning of the grip (Fig. 4.7A row 3 and B). This hand specific bilateral pattern becomes more ipsilateral after 9 s of hand-gripping. This kind of ipsilateral activation has been previously observed in strenuous force generation but only after more extended stimulation periods ($\geqslant 30\,\mathrm{s}$) (Shibuya et al., 2008, 2016). HemCNN allows us to find similar patterns just after 9 s.

HemCNN supports bilateral and hand-specific hemispherical activations Our results can also be interpreted in the broader extent of motor control. The seminal cortical electrostimulation study of Penfield and Boldrey (1937) established that the contralateral cortex was directly connected to the opposite upper limb. This direct connection resulted in the opposite upper limb's elementary movements when the contralateral motor cortex was stimulated, which did not happen when stimulating the ipsilateral motor cortex. Since then, finding contralateral brain activation patterns is customary in motor control studies. However, these findings did not rule out the possibility of a simultaneous and ipsilateral activation modulating the contralateral activity during more elaborated movements. Indeed, virtual lesions in the ipsilateral motor cortex lead to loss of motor performance in the contralateral hand (Davare et al., 2006). The current opinion gears towards the idea that both hemispheres act together to produce complex unimanual and bimanual movements (Serrien et al., 2006; Haar et al., 2017; Ames and Churchland, 2019). Notably, in hand force evoked HRs analysed with fMRI, the lateralisation is often reflected as a dominant activation during the task rather than an active versus non-active hemisphere (Ehrsson et al., 2000; Keisker et al., 2010).

In contrast, for fNIRS identifying lateralisation in the HR has been challenging (Shibuya et al., 2008; Derosière et al., 2014; Shibuya et al., 2016; Wriessnegger et al., 2017). The literature points at the effects of systemic artefacts (Tachtsidis and Scholkmann, 2016; Caldwell et al., 2016; Pfeifer et al., 2018) and to differences in the motor tasks as reasons for the lack of specific hemispheric patterns. In the first place, the extent of muscular mass involved in the task seems to have a role. More significant

amounts of muscular fibres recruited in tasks involving whole hand grips seem to evoke more global responses (Shibuya et al., 2008; Nambu et al., 2009; Holper et al., 2009; Derosière et al., 2014; Shibuya et al., 2016) compared to tasks recruiting a reduced amount of muscles in finger pinches (Nambu et al., 2009) or no muscular mass at all in force imageries (Coyle et al., 2004; Sitaram et al., 2007) that evoke more contralateral responses. In terms of the HR time-resolution during hand grips, a 180 s long force-ramp grip task allowed Shibuya et al. (2008) to find an early (100 – 180 s) contralateral and a delayed (130 – 180 s) ipsilateral activation. In contrast, more recently, Derosière et al. (2014) could not identify such timed lateralisation differences during a shorter (30 seconds) steady grip.

We argue that the lack of such timed lateralisation patterns in the previous studies, with slower hand contractions than ours, is not only a consequence of systemic artefacts as previously suggested (Pfeifer et al., 2018). It is also a consequence of using linear or quasi-linear methods that find differences using steady or fixed descriptors of the hemodynamic response (e.g. amplitude or canonical responses, respectively). The drawback is that these descriptors are likely to be modulated by the task (Leff et al., 2011), but conventional methods fix them instead of adapting them to the actual HR shape (Abdelnour and Huppert, 2009). We show how hand-specific hemispheric temporal patterns can be resolved for lateralised and fast hand grips using a data-driven non-linear method. Our hemCNN method simultaneously optimises the localisation of hemodynamic response transients, their time alignment, and their channel relationship.

Limitations We demonstrate that hemCNN is particularly suited for interpretation, which other CNN approaches fail to provide (Hiwa et al., 2016; Trakoolwilaiwan et al., 2017; Ghonchi et al., 2020). In particular, hemCNN captures functional HbO/HbR signatures and interhemispheric differences. However, we note that our approach's implementation is limited to binary classification tasks aiming at finding such interhemispheric differences. In hemCNN, each output unit is exclusively informed by each hemisphere in a comparative way, resulting in only two possible outputs. Since each output represents a different class, only two classes drive the differentiation of hemispheric activities. If one wanted to find hemispheric differences across a greater number of classes, modifications in the architecture would be required. A potential modification would be to use as many hemCNNs as classes we need. Each hemCNN would be trained at the same time to predict the absence or presence of a particular class. Nonetheless, the way in which the dropout is designed should reflect in the presence of several classes and adapted adequately, so there is a balanced reflection of each class on each hemisphere. Additionally, we cannot discard that the visual stimuli might have been driving the response. However, we consider

this unlikely given the hand-specific topographical around C3 and C4.

Throughout the present work, we showed how the potential power of generic CNNs was tailored to particular research hypotheses with interpretability in mind. We expect that our work will facilitate the analysis of fNIRS aiming at finding hemispherical roles in the performance of complex tasks. It also has the potential to inspire further exploration of interpretable artificial neural networks methods.

4.5 Conclusion

We introduced an interpretable CNN based method to temporally and spatially resolve cortical Haemodynamic Responses (HRs) during a motor task that evokes highly variable ones. Our results can be interpreted in the light of current fNIRS results for similar but less dynamic motor tasks. For a neuroimaging technique that delivers hemodynamic responses modulated by the task frequency, complexity and invested effort and that are affected by systemic artefacts, the hemCNN method offers new analytical and decoding opportunities. Our results suggest that lateralisation features could be present in the cortex for all sorts of motor tasks but that the complexity of the HR hinders the conventional methods' ability (often linear) to resolve it. In particular, hemCNN can be directly applied in real-time analysis since the convolutional operator is invariant to time translations, differences are found across time-samples and channels, and there is no baseline correction or strong preprocessing other than the z-scoring of the fNIRS example. This makes hemCNN appropriate not only for offline neuroimaging but also for BCI paradigms. In the first case, it can be useful for group studies where models of the HR are not accurate due to abnormal function, such as in stroke. In the second, it can be beneficial for rehabilitation or assistive BCIs where the response of the BCI needs to be decoded from the users' attempts of producing force.

Chapter 5

HYGRiiP: Multimodal full brain-muscle characterisation of bimanual grip[1]

5.1 Introduction

Brain-Computer Interfaces (BCIs) are communication systems that bypass the natural nerve-muscle communication pathway. They directly exploit brain signals to communicate with the environment through a mechanical device. This kind of interfaces offer new opportunities for people with disabilities to assist (Onose et al., 2012), rehabilitate (Daly and Wolpaw, 2008; Soekadar et al., 2011; Enriquez-Geppert et al., 2013; Ang et al., 2015) and even restore (Bouton et al., 2016; Ganzer et al., 2020) their motor control. This is particularly important for the hands, which we constantly use to interact with our surroundings. Among people with tetraplegia, hand function is the most desired motor ability to recover (Anderson, 2004; French et al., 2010). Additionally, after stroke, targeting hand force during rehabilitation leads to longstanding improvement of hand function (Canning et al., 1999; Ada et al., 2006).

Indeed, force generation is essential in our physical interactions (Westling and Johansson, 1984; Ostry and Feldman, 2003). From a functional perspective, force enables us to stabilise the dynamic loads of our naturally unstable interactions with the environment and tools (Franklin et al., 2003; Faisal et al., 2010). We also need to regulate our force output to manipulate fragile and plastic materials. Even from an aesthetic point of view, force allows us to imprint character when playing musical instruments.

[1] Ortega, P., Zhao, T., & Faisal, A. A. HYGRiiP: Full-stack characterisation of neurobehavioural signals during a bimanual grip force tracking task. In review with *Nature Scientific Data*.

Moreover, from a social perspective, we enrich our non-verbal expression by applying different degrees of force, for example, during a hand-shake. Additionally, the brain signals controlling the modulation of force output are naturally occurring in the cortex (Cheney and Fetz, 1980). Understanding how these brain signals modulate the force output might open new venues for BCI continuous control.

As a result, the study of force has received much attention in BCI. Especially, the invasive BCI field has achieved important milestones in understanding (Phillips, 1977; Cheney and Fetz, 1980; Porter and Lemon, 1993), decoding (Carmena et al., 2003; Chhatbar and Francis, 2013) and even restoring (Bouton et al., 2016; Ganzer et al., 2020) the generation of force. However, their application is still limited by the need to undergo surgeries, and many of their target users cannot afford surgery due to its cost or health complications.

In contrast, non-invasive approaches —a completely safe and cheaper alternative— have only made moderate progress in force decoding (Nambu et al., 2009; Paek et al., 2019; Hayashi et al., 2019). One of the reasons explaining this is the lack of understanding of how non-invasive brain signals modulate force (Wriessnegger et al., 2017; Paek et al., 2019).

In this sense, data-driven approaches might offer new opportunities in finding the non-invasive brain signal features that improve force decoding. We have already shown how hemCNN, a CNN approach for fNIRS processing (Chapter 4), enables the analysis and decoding of haemodynamic cortical responses during force generation. In particular, we showed that an appropriately designed data-driven technique unlocks fNIRS to be used to study strenuous force generation, where conventional linear approaches failed. In addition to fNIRS, previous research has also shown the advantages of combining fNIRS with EEG to improve the decoding of motor variables (Hansen et al., 2014; Yin et al., 2015; Chiarelli et al., 2018; Ghonchi et al., 2020).

However, non-invasive multimodal research has yet to explore the continuous generation of force. In part, this is due to the lack of suitable multimodal decoding methods (Yin et al., 2015). Moreover, most non-invasive studies only focus on the unilateral motor control, i.e. control exerted with either the right or left hand (e.g. Nambu et al., 2009; Wriessnegger et al., 2017; Paek et al., 2019; Hayashi et al., 2019). To the best of our knowledge, the study of force control with both hands simultaneously during dynamic force generation is yet to be investigated using non-invasive multimodal approaches.

Here, we present HYGRiiP: A ready to use dataset for the study and decoding of bimanual dynamic force generation (summary of characteristics in Table 5.1). HYGRiiP includes multimodal brain signals

dataset	N	DWT [h]	trials	task	conditions	available at
HYGRIIP	9	9.4	1350	bimanual	5 contraction targets	data/repo
HYGRIP	14	2.8	328	L/R unimanual	2 (left/right hand)	data/repo

Table 5.1: HYGRiiP and HYGRiP short comparative. N stands for "number of subjects" and DWT [h] for the "data-worth time in hours", i.e. the total amount of hours that contain exploitable data. In HYGRiiP, both hands are used simultaneously to generate dynamic force, introducing the exciting problem of 2 simultaneous DoF control in non-invasive BCI. As a reminder, HYGRiP consists of a unimanual task where the left/right (L/R) hands are used in different trials. The "i" character in HYGRiP and HYGRiiP and their logos represents the human arm and visually conveys the simultaneous use of one or two hands, respectively. The repositories digest the raw data into a format that makes it suitable for immediate use by machine learning engineers (Lawrence, 2017) who do not need to go through a lengthy cleanup and reshaping process.

in the form of co-aligned EEG and fNIRS signals. We also include continuous motor behaviour signals using bilateral forearm EMG and bimanual isometric force signals. Especifically, the task used in HYGRiiP requires both hands to be used simultaneously. We design hand specific target profiles that participants need to track with each hand at the same time. In particular, we make both hands to be active or relaxed at the same time. It is only during the contraction that each hand follows a different force trajectory. In this way, we aimed to evoke hand-related cortical activations only as a result of the dynamic contraction generated by each hand.

HYGRiiP opens up an interesting question in the non-invasive BCI field. Namely, with this dataset, we interrogate the possibility to control 2-DoF simultaneously, i.e. one per hand. To date, non-invasive BCI only achieves different degrees of freedom by sequentially controlling each of them via discrete EEG events (Wolpaw and McFarland, 2004; Royer et al., 2010; Meng et al., 2016). This means that a user needs to wait to finish controlling one degree of freedom to act on the next. Consequently, the time that a user needs to complete a task is much higher than if the user could operate both DoF at the same time, which is the case introduced in this new dataset.

With HYGRiiP, we aim to inspire and stimulate the development of advanced multimodal BCI techniques for the decoding and study of bimanual dynamic force generation. This will facilitate the development of suitable machine learning techniques for the exploitation of multimodal measures. We believe this will advance the use of multimodal BCI with the ultimate goal of improving the autonomy of people with disabilities.

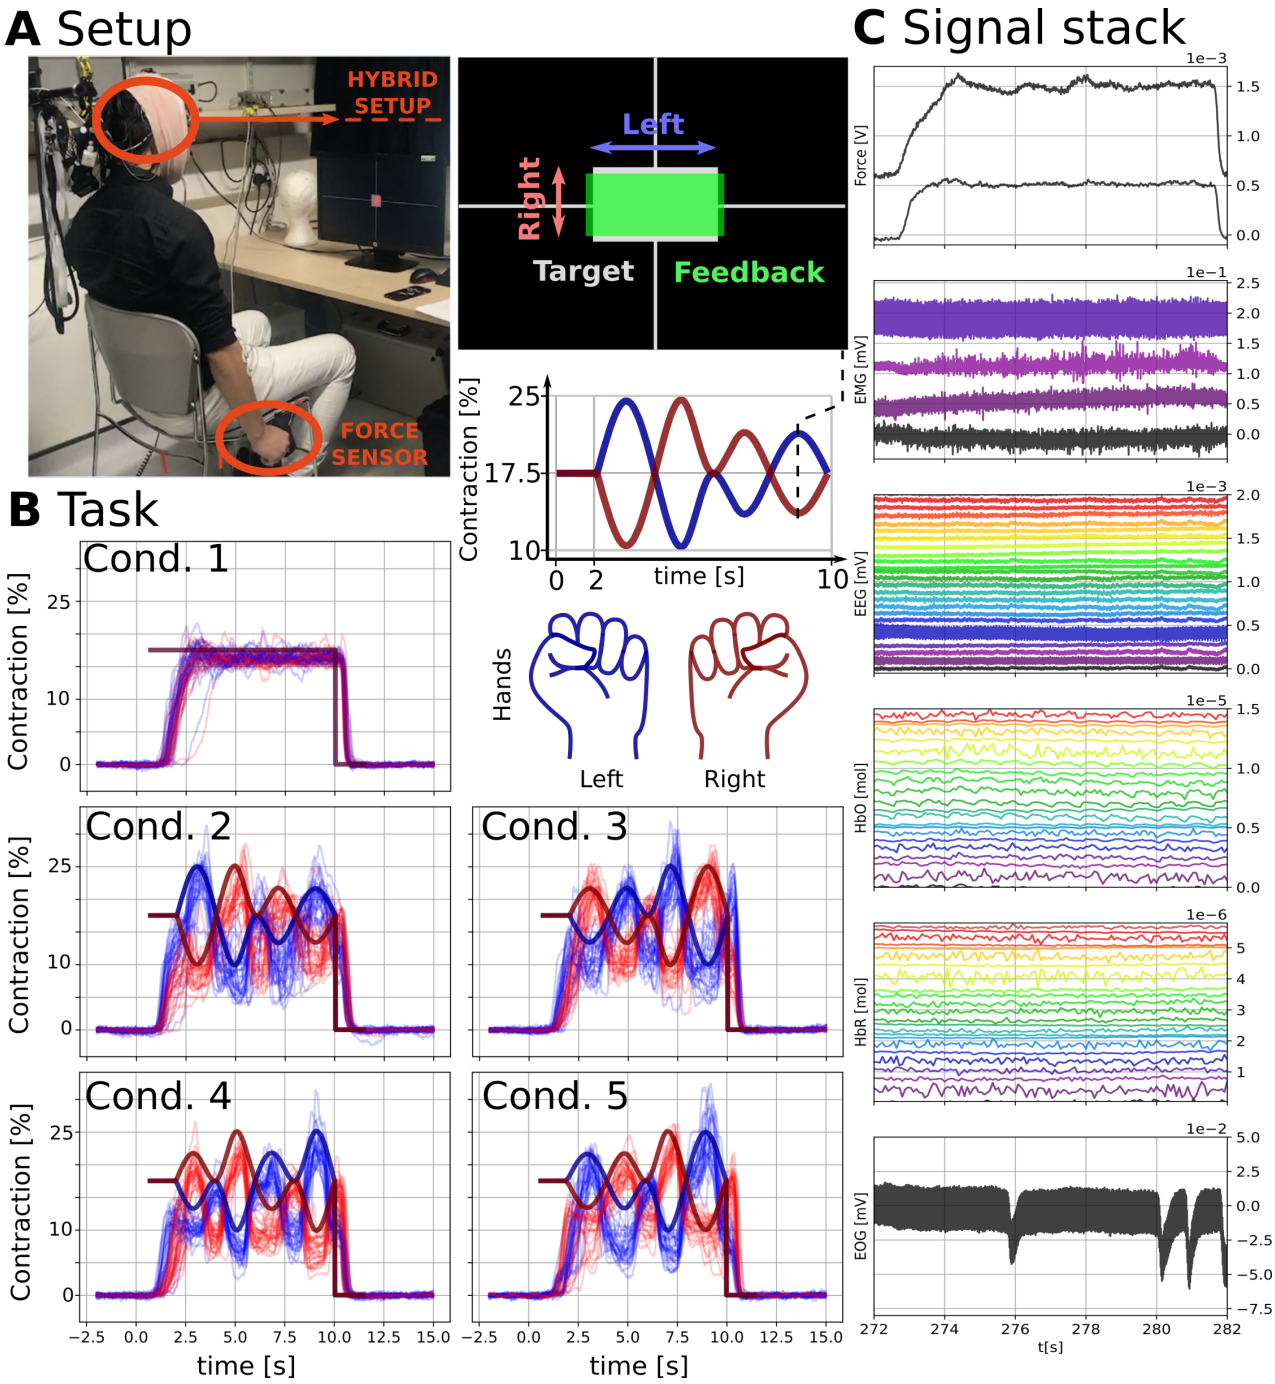


Figure 5.1: **A** (Left) Experimental setup with visual feedback. (Right) Schematics of the task contraction targets and visual feedback. The grey rectangle width indicates the force to be applied with the left hand and the height, the force to be applied with the right hand. Below, a schematic example of the task with force trajectories for the left (blue) and the right (red) hands. Maximum contraction is set to 25% MVC and the minimum to 10% MVC. **B** The layout of the 5 conditions used for the experiment. The task is presented with thicker and darker colours overlaying the executed forces. The presented behavioural results belong to subject Q. **C** A subject's full-stack of neurobehavioural data (10 s, selected channels per modality for readability) as coming out of the respective devices, i.e. not preprocessed. (Top-bottom) Force produced by subject's right and left hands; EMG channels from right and left forearm; EEG channels; Oxygenated Haemoglobin (HbO); Deoxygenated Haemoglobin (HbR); EOG with 4 eye blinks. Units in the bottom x-axis and each corresponding y-axis.

5.2 Methods and materials

5.2.1 Participants

Nine (N = 9, anonymised IDs from "O" to "X") healthy, right-handed volunteers participated in this dataset production. As for HYGRiP, handedness was confirmed by the Edinburgh inventory (Oldfield, 1971) for all participants. No participant reported a history of neurological, cardio-respiratory or physical disorders. Imperial College Research Ethics Committee approved all procedures, and all participants gave their written informed consent. The experiment complied with the Declaration of Helsinki for human experimentation and national and applicable international data protection rules.

5.2.2 Motor control task

In HYGRiiP, the motor task consists of a bimanual isometric hand-grip with 5 different conditions. Each condition consisted of 10 s duration trials where the participants had to track a force with each hand. In condition 1, both hands had to produce the same flat level of contraction for the trial duration. In the remaining 2nd to 5th conditions, each hand had to track a variable contraction trajectory as shown in Figure 5.1.

Participants sat down in front of the computer screen, delivering visual feedback. Their arms remained naturally hanging straight down, and force-grip sensors were placed within their immediate reach on both hands. Before the start of the experiment, the Maximum Voluntary Contraction (MVC) was recorded and used to calibrate each hand's contraction to deliver the contraction targets of each task.

The level of contraction was guided by a grey square placed at the centre of the screen. The square's width defined the contraction that needed to be produced by the left hand and the height that to be produced with the right hand. The actual levels of contraction produced by the participants were fed back on the screen using an equivalent square with the actual contraction levels instead of the target ones. The feedback rectangle was set with 0.5 transparency, allowing the subjects to see the target rectangle. This rectangle was colour coded, so it turned green when subjects stayed in the ±10% error range from the target and red otherwise (Fig. 5.1A left). The level of contraction during condition 1 was set to 17.5% MVC. For conditions 2-5, the highest level of contraction was set to 25% MVC and the lowest to 10% MVC. Both hands' trajectories were designed to have an average 17.5% MVC contraction, as in condition 1. The trial's start was indicated with a "beep", and for the

first 2 s of conditions 2-5, the target indicated a flat 17.5% MVC that the subjects aimed to reach before the dynamic part of the task started. All participants did a balanced amount of 30 trials per condition, except two participants who were too tired to perform the last block and requested to stop the experiment. Conditions were pseudorandomised across trials to avoid anticipation and interference between them. Each trial was followed by a randomised resting period uniformly distributed between 15 and 21 seconds, to avoid phasic constructive interference of systemic artefacts, e.g. Mayer waves or breathing, in the brain responses. The refreshing of the feedback in the screen was set to 100 Hz.

5.2.3 Data collection pipeline & Methods

We used the same setup and recorded the same signals that we did for HYGRiP (cf. Chapter 3) with some exceptions. First, we recorded an additional force channel using another grip force transducer (PowerLab 4/25T, ADInstruments, Castle Hill, Australia) to allow for the left and right-hand grips to be recorded simultaneously. Second, instead of recording breathing activity, we recorded the skin hemodynamic activity on the forehead.

The skin hemodynamic activity in the form of HbO concentration changes was recorded using a NONIN 8000R (NONIN, Tilburg, The Netherlands) placed on the forehead sampling at 75 Hz. The forehead was selected as the measurement location for being the closest portion of skin to the scalp with better access. Nonetheless, the presence of several recording systems in the surrounding area (fNIRS, EEG and EOG) limited the space on the forehead, which resulted in unstable skin hemodynamic activity measures for some subjects and sessions. In total, 17 sessions were stably recorded and included in the dataset. Finally, some of the sampling frequencies were modified. EEG, EMG and EOG were resampled at 1000 Hz and forces at 200 Hz. The rest of the measures were sampled with the same frequencies of HYGRiP.

5.2.4 Dataset overview

We provide the dataset as a single hard-disk file (HDF) which has undergone very little processing to avoid biasing future analyses. The companion repository (Table 5.1) can be used to process the signals with ease.

The dataset file structure follows a tree-like organisation in three levels: `subject/condition/measure`. At the `subject/condition` level, the attributes now include more details on the task and the recorded

feedback provided at the screen for every frame. For example,

```
Subject P:
    Measure P/1:   date   events    exec_order    screen_center
                   screen_counts    screen_frames screen_xpos   screen_ypos
    Measure P/1/bhv:   events
    Measure P/1/dxy:   events
    Measure P/1/eeg:   events
    Measure P/1/emg:   events
    Measure P/1/eog:   events
    Measure P/1/frc:   MVC   events
    Measure P/1/oi1:   events
    Measure P/1/oi2:   events
    Measure P/1/oxy:   events
```

5.3 Data validation

5.3.1 Data validation preprocessing pipeline

The following preprocessing was applied to the raw signals in the dataset to obtain the brief analysis in Figures 5.1 to 5.4 which can be reproduced using the companion notebook. After preprocessing all signals, epochs were extracted from 2 s before the "Go" instruction to 15 s after.

fNIRS. Optical intensities were low-pass filtered below 0.25 Hz with a 7th order elliptical filter. Changes in haemoglobin concentration were obtained using Eq. 3.2 (see Chapter 3). Figure 5.2B shows the average of the HR across all subjects and conditions. As in HYGRiP, we can see an apparent increase of HbO and decrease, at a smaller scale, of HbR associated with the onset and end of the force at 0 and 10 s of the task.

EEG. EEG was first downsampled to 250 Hz (with anti-aliasing down-pass filtering). Notch filters were applied at the mains (50 Hz) and fNIRS (12.5 Hz) frequencies and their harmonics. EEG was finally high-pass filtered above 1 Hz using a 5th order Butterworth filter. ICA was then applied to automatically remove EOG artefacts, rejecting 1 component when the independent EEG component and the EOG had a correlation above 0.3 in absolute value.

Figure 5.2A shows the averaged spectrogram of the signal in the mu-band (9 – 12 Hz). As we can see, the most obvious feature is the mu-power suppression during movement and its rebound at the

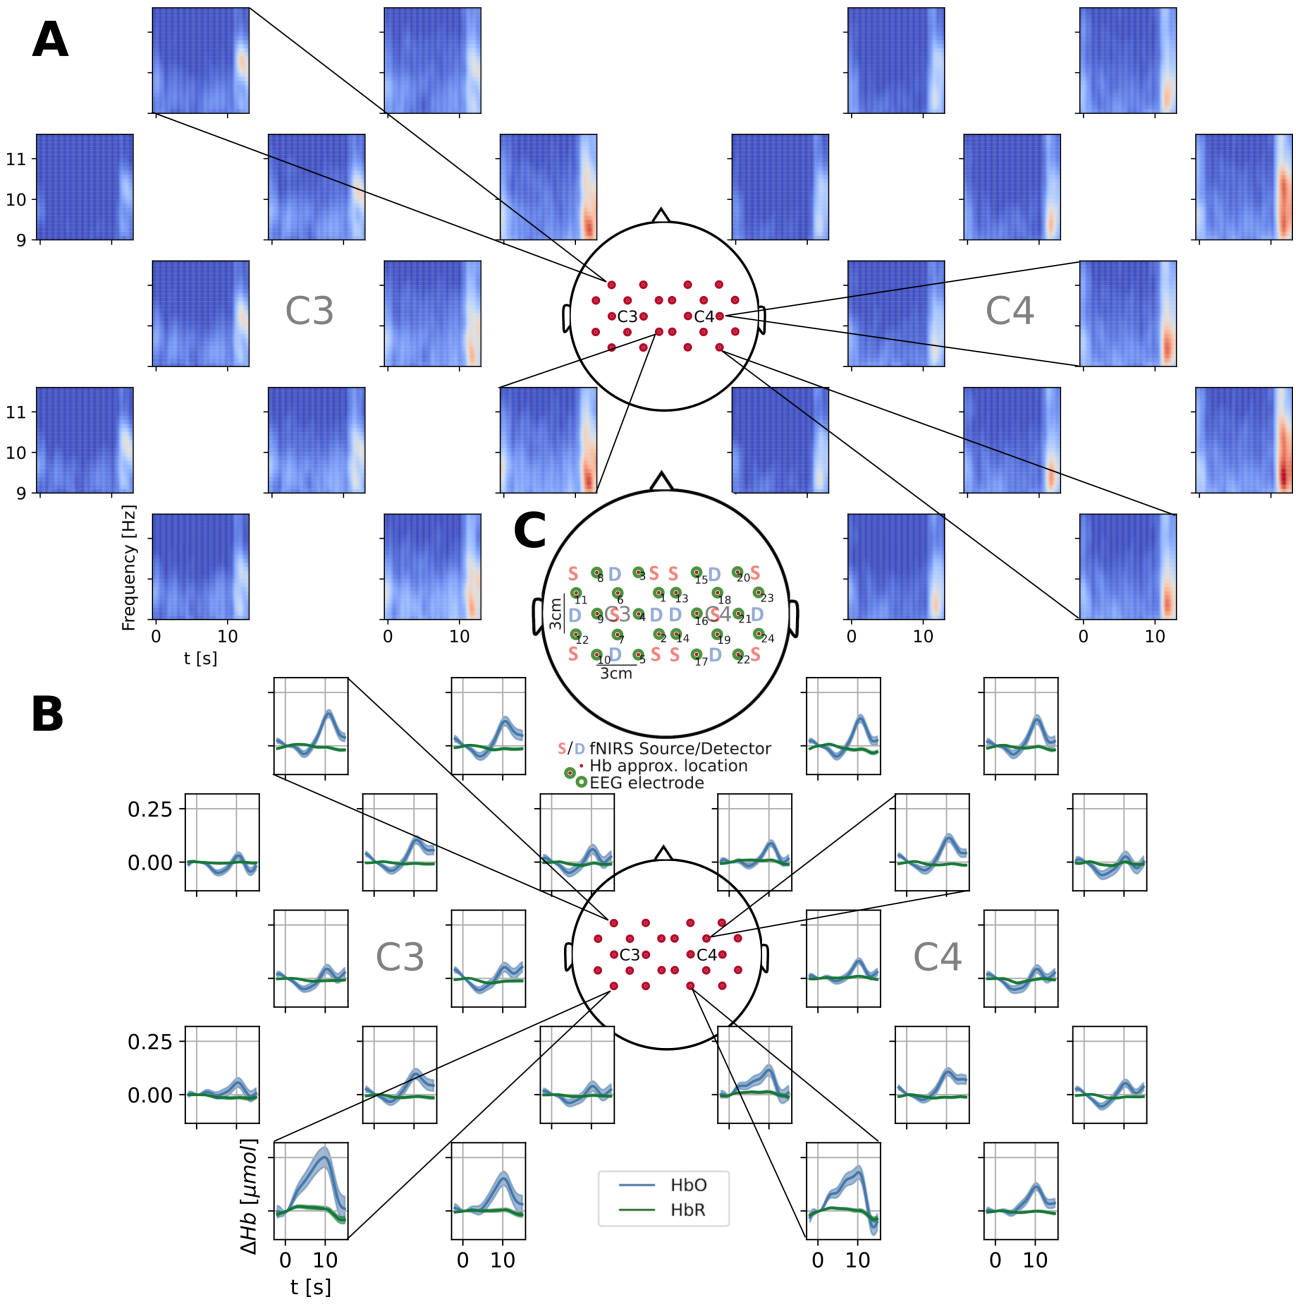


Figure 5.2: **A** Cross-subject mu-band spectrogram averaged across right-hand trials time-locked to the 'Go' cue ($t = 0\,\mathrm{s}$) of each trial. All conditions are averaged, giving subtle bumps of power changes since each condition has a different force trajectory to track and force production is not temporally aligned. The most obvious component is the power rebound once force stops being applied. Something that happens at the same time in all conditions. **B** Cross-subject average of HbO and HbR changes across trials. HbO increases during the task, and HbR decreases at a smaller scale. Both start returning to baseline after the 10 s of activity. Units are shown in the bottom-left corner. **C** Channel layout. The numbers correspond to the row position in the dataset arrays.

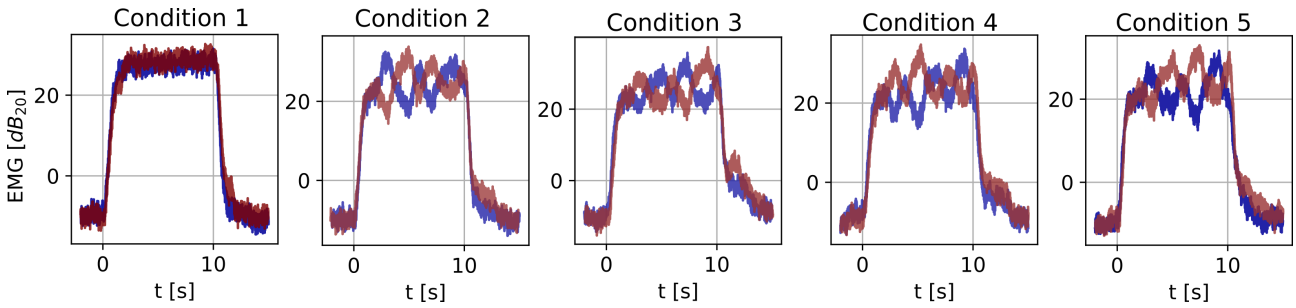


Figure 5.3: Cross subject average of the Hilbert power envelope of the EMG for each condition. On average, the power envelope follows the trajectory of the target forces per hand and condition. As usual, red corresponds to the right hand and blue to the left.

end of the task (Cochin et al., 1999; McFarland et al., 2000). In particular, the task has one force onset followed by a continuous gradation of force that never reaches zero, i.e. relaxation, until the task finishes and the subject relaxes. Therefore, the re-synchronisation of the neural population is more evident in the end when the muscle completely relaxes. Nonetheless, traces of periodic increases and decreases of power can be identified as lighter blue bands during the task (Fig. 5.2A). Note that we averaged across all conditions. The force trajectories across conditions do not have aligned increases and decreases of force. This might be reducing the changes of mu-band power related to the task when we average across conditions.

EMG. The EMG was also first downsampled to 250 Hz and then high-pass filtered above 110 Hz with a Butterworth filter of order 17th. Figure 5.3 shows the power envelope of the EMG per condition.

Force. The force signal was first resampled to 250 Hz, then band-pass filtered between 0.01 and 10 Hz with a Butterworth filter of order 3 and then again high-pass filtered with an elliptical filter of order 1 above 0.01 Hz. Drift was further reduced by removing the linear drift per trial. The force measures were finally converted to contraction values using the recorded MVCs.

EOG. EOG was downsampled to 250 Hz and then high pass filtered above 1 Hz.

5.3.2 Correlations across multimodal signals & validation

Signals were correlated to characterise a "brain and behaviour" structure (EEG, fNIRS, EMG and force in Fig. 5.4). We computed the correlation in the 2 s window around the "Go" signal (Fig. 5.4A) and

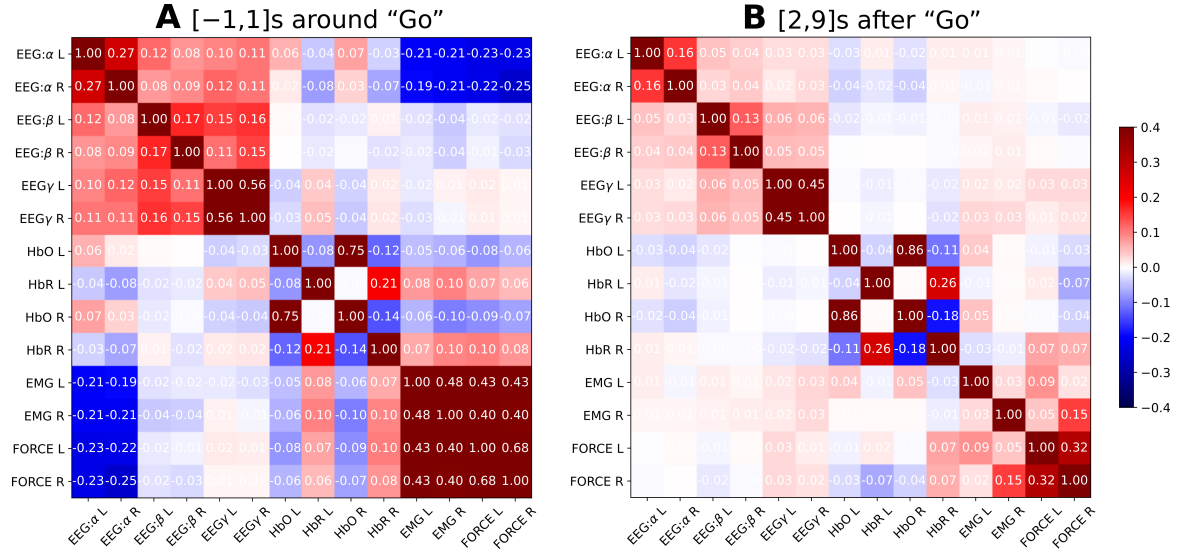


Figure 5.4: Correlation between brain and behavioural measures at difference trial times. Only correlations with $p < 0.001$ are shown. On the left, the correlation is computed in the $[-1, 1]$ s window around the “Go” of the trial. We can observe a similar correlation structure as in HYGRiP, except that now both hands present high correlation as a result of being used at the same time. On the right, the correlation is only computed during the task $[2, 9]$ s after the “Go”. As we observed earlier for HYGRiP, the correlation between the EEG, the force and the EMG is much lower in this interval than at the beginning of the task. In this case, there is a continuous contraction that only varies in intensity which the phase of ERDs might not be able to capture due to the continuous nature of the desynchronisation.

can observe a similar ERD (Pfurtscheller and Lopes Da Silva, 1999; Pfurtscheller et al., 2006) in the mu-band with respect to the initiation of the activity of the EMG and the force.

Compared to HYGRiP (Chapter 3), here, both hands are used at the same time. Consequently, we now observe the two forces and EMG channels negatively correlating with the mu-band. As advanced earlier for HYGRiP, and now showing in Fig. 5.4B, this correlation diminishes during the dynamic part of the trajectory possibly due to the lack of resolution of the phase change of the ERDs.

We also observe differences in the cross-hand correlation of force and EMG. For the force onset, i.e. the part of the task when subjects initiate the contraction from a relaxed state (Fig. 5.4A), the correlation of left and right hand is around $r \approx 0.5$ for EMG and $r \approx 0.7$ for the force. For the task's part restricted to the dynamic trajectory (Fig. 5.4B) the same cross-hand correlation falls to $r = 0.03$ and $r = 0.32$ respectively, with most of that correlation being due to condition 1. This correlation illustrates how the task can be separated qualitatively into a “force onset” and “dynamic force tracking” since the correlation structure of the behaviour changes. In the first case, during the “force onset”, both hands start producing force simultaneously. In the second case, during “dynamic force tracking” (except for condition 1), the forces are produced with low correlation and are more independent.

Finally, we corroborate that the skin hemodynamics reflecting the breathing and pulse fluctuations of blood oxygenation do not have force prediction power (Fig. 5.5). For the 17 available sessions, we

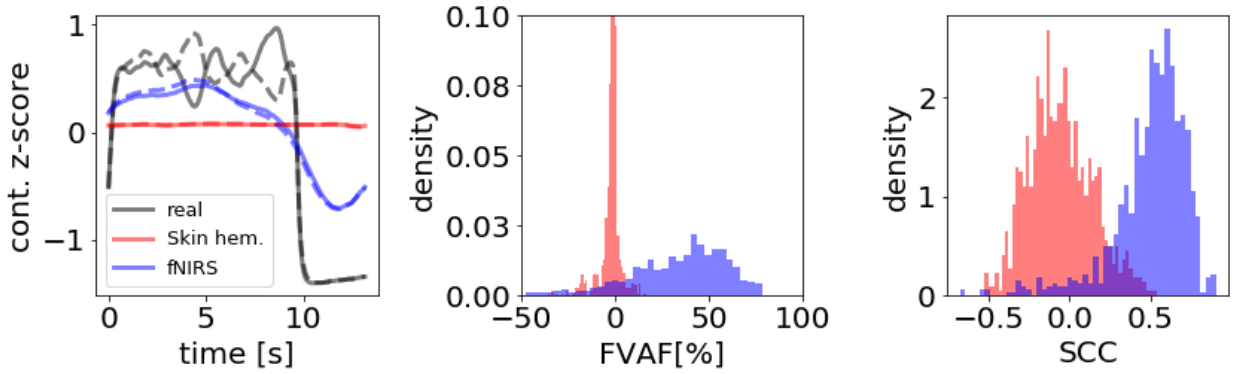


Figure 5.5: Scalp hemodynamics have negligible force-predicting power. (Left) Z-score of the average contraction (black) for the left (discontinuous) and right (continuous) hands. Predicted z-score of average contraction from fNIRS in blue and from scalp skin haemodynamics in red. The contraction decoded from fNIRS signals can capture the increase and decrease of force while the contraction decoded from skin haemodynamics is flat. (Centre) Distribution of the Fraction of Variance Accounted For (FVAF) for all trials. FVAF measures the per cent reconstruction of force trajectories. When using fNIRS (blue), we achieve around a 50% reconstruction of the original force. If skin haemodynamics are used, the reconstruction is around 0%. (Right) The same can be observed using Spearman's Correlation Coefficient (SCC) as a reconstruction metric.

built a linear model decoding the force from the scalp hemodynamics and another one using fNIRS. We then computed two metrics, FVAF[%] and SCC, that measure the original signal's reconstruction. The fraction of variance accounted for (FVAF[%]) represents the original signal's percentage that the prediction can reconstruct. A 0% FVAF represents a 0% reconstruction while a 100% means the signal can be reconstructed entirely. Spearman's correlation coefficient (SCC) is a robust measure that computes signals' correlation when these do not have normal distributions. An SCC close to zero means that the signals are not correlated. Values closer to 1 or −1 mean that the signals are correlated or anti-correlated, respectively.

The results indicate that skin hemodynamics have a negligible prediction of force (SCC = −0.07 and FVAF[%] = −1%) compared to using fNIRS (SCC = 0.5 and FVAF[%] = 32.5%); prediction $[0, 14]$ s of force profile, Wilcoxon test, $p < 0.001$, Fig. 5.5). In particular, we can observe how the average predicted contraction from fNIRS rises after the force onset and decays after the contraction ends. In contrast, the contraction reconstructed by skin hemodynamics remains flat.

Note that the 17 available sessions have been used to train and test. Hence, the results here should not be taken as a predictive measure of how well fNIRS can reconstruct force. Instead, these results reflect the poor force reconstruction capabilities of skin haemodynamics even when overfitting is allowed. In summary, systemic variations of haemoglobin concentrations present in fNIRS are unlikely to drive the decoding of force.

5.4 Conclusion

We have presented here HYGRiiP, an exhaustively collected, compiled and validated multimodal brain and behaviour signals dataset. In particular, HYGRiiP opens up the challenge for 2-DoF simultaneous BCI control. To date, non-invasive BCI achieves the control of several DoF by operating each DoF separately and in sequence, increasing the time needed to complete simple daily tasks. The control of several DoF is vital for hand forces since many daily tasks involve the bimanual application of forces. In these circumstances, simultaneously operating 2DoF would significantly reduce the time needed to complete daily tasks. However, algorithms suitable for the fusion of multimodal brain signals (EEG and fNIRS) for continuous control are still lacking.

Our task reflects on the simultaneous and continuous control of both hands reflecting into real-world manipulations (Lima et al., 2020; Haar and Faisal, 2020). The ready to exploit data in HYGRiiP is intended to stimulate the development of these novel machine learning algorithms and advance the field of multimodal non-invasive BCI algorithms. With HYGRiiP, researchers can seamlessly start developing their algorithms and analyses straightaway. Our ultimate goal is to help in the collective effort to deliver a safer and more affordable BCI technology to improve the autonomy of people with disability.

Acknowledgements

We acknowledge the collaboration of Tong Zhao in coding the task, recruiting participants and recording the data.

Chapter 6

Decoding bimanual force from multimodal (fNIRS + EEG) non-invasive brain signals with Deep Learning[1]

6.1 Introduction

Brain-Computer Interfaces (BCIs) offer an alternative pathway to interact with our environment. These interfaces are especially relevant for people whose natural ability to physically interact with their environment has been lost or damaged (Wolpaw et al., 2000; French et al., 2010; Branco et al., 2019b).

During these mechanical interactions, the decoding of forces is essential for stability, safety, and imprinting character in the mechanical interaction (Xiloyannis et al., 2017; Branco et al., 2019a). Research in invasive BCIs has also shown that force decoding can provide more generalisable BCI decoders (Cherian et al., 2011), specially in changing dynamic conditions (Chhatbar and Francis, 2013). Additionally, understanding force decoding might offer more intuitive control of BCIs since the signals already exist in the brain (Flint et al., 2014).

These force decoding benefits in BCI acquire particular importance for the hands. As the specie with

[1]Published in: Ortega, P. & Faisal, A. A. (2021). Deep real-time decoding of bimanual grip force from multimodal EEG & fNIRS sensor fusion. In *2021 10th International IEEE/EMBS Conference on Neural Engineering*. IEEE. Submitted as journal article: Ortega, P. & Faisal, A. A. Deep Learning multimodal fNIRS & EEG signals for bimanual grip force decoding. Submitted to the *Journal of Neural Engineering*.

the greatest dexterity, we spend most of our day interacting with our surroundings using our hands. Indeed, for completely paralysed people, the hand function's recovery is the most relevant priority (Anderson, 2004; Snoek et al., 2004; French et al., 2010) and hand force recovery is one of the first rehabilitation targets after stroke (Ward et al., 2003). Hand force generation needs have been vastly met in invasive BCIs (Carmena et al., 2003; Bansal et al., 2012; Flint et al., 2014; Naufel et al., 2019). However, these invasive approaches do not reach most of the potential BCI users. In particular, invasive BCIs require severe surgical procedures that many users are not eligible for or cannot afford. In contrast, non-invasive BCIs offer a plug-and-play cheaper alternative with zero impact on the users' health. Nonetheless, compared to invasive approaches, non-invasive BCIs have explored human force decoding to a lesser extent.

Discrete force characteristics have been decoded from multimodal (EEG and fNIRS) signals. Hansen et al. (2014) decoded two force levels, 20% and 60% of Maximum Voluntary Contraction (MVC), and detected force produced with the right foot. Their multimodal approach achieved a 76% accuracy in foot force detection and an 81% accuracy in force level classification. They also found that the use of only one modality led to a drop of approximately 10 points of accuracy in force detection with fNIRS, and a similar decrease for force level classification with EEG. This suggests that fNIRS and EEG provide complementary information with fNIRS better at decoding force levels and EEG at detecting force generation. Similar multimodal approaches have been used to classify force imageries showing increase in classification accuracy compared to the unimodal systems (Yin et al., 2015; Fu et al., 2016). Altogether, these studies showed that the combined use of fNIRS and EEG provides a significant advantage in the classification of discrete force characteristics.

Nonetheless, all these approaches based their classification using complete trials instead of classifying these characteristics each time-step, which would enable a continuous control of a device (Paek et al., 2019). Indeed, despite the benefits introduced by these multimodal approaches, EEG or fNIRS have only been separately used to decode continuous force. Nambu et al. (2009) used a linear technique to regress fNIRS signals to finger pinch forces achieving a FVAF[%] $\approx 45\%$ reconstruction of a smoothed version of the produced force. However, their approach cannot be strictly considered decoding since it was not causal and used information in the future for the regression. With EEG, Paek et al. (2019) decoded hand-grip force trajectories and also reported a modest reconstruction performance (correlation ≈ 0.42). Finally, Hayashi et al. (2019) continuously decoded a three discrete level force task achieving an adjusted R^2 of 0.62. Nonetheless, their method required the prior knowledge of

which two out of the three force levels a user needed to differentiate. Ultimately, they did not show a way of decoding the three combined levels in a continuous way.

Overall, multimodal or unimodal non-invasive BCIs have only used linear methods to decode force (Hansen et al., 2014; Fu et al., 2016; Paek et al., 2019; Hayashi et al., 2019). In contrast, invasive studies suggest that the relationship between the brain signals and the exerted force can be better captured with non-linear models (Naufel et al., 2019; Fatemi and Daliri, 2020). In this context, Deep Learning (DL) (LeCun et al., 2015) –a flexible machine learning framework for non-linear modelling– can introduce further advantages in non-invasive BCI force decoding. It has been shown that DL is a powerful tool for EEG classification (Schirrmeister et al., 2017; Ortega et al., 2018). Furthermore, in Chapter 4 we also showed how DL outperforms standard fNIRS processing methods by automatically finding the best spatio-temporal features of the signal for a force task. In addition to the lack of exploration of non-linear techniques, non-invasive BCIs have been limited to the decoding of unilateral hand activities. In contrast, bimanual interactions are more valuable in daily activities.

In this chapter we explore bimanual continuous force reconstruction from multimodal (fNIRS and EEG) non-invasive cortical signals with DL. We first ask if multimodal signals (fNIRS and EEG) provide an advantage in bimanual continuous force decoding compared to each separate modality. For comparison, we used Paek et al. (2019) linear approach to do this. Secondly, we explore DL techniques to fuse these multimodal signals with different neurophysiological origins (Ball et al., 1999).

Our bimanual task allows us to specifically answer questions about bimanual force decoding in the continuous case. Namely, we evaluate three decoding performance aspects: (1) **force trajectory reconstruction**, i.e. how well the methods can reconstruct bimanual force trajectories from non-invasive signals; (2) **force detection performance**, i.e. how well models can detect if the hands are applying force at each time-step; and (3) **hand disentanglement**, i.e. how accurate are the models at distinguishing which hand is producing more contracted, i.e. is more active, at a given time.

In addition to analysing the decoding capabilities, exciting questions are introduced by using a non-linear decoding method. Previous studies suggested that non-invasive brain signals are not linearly modulated with force (Paek et al., 2019; Hayashi et al., 2019). We hypothesise that the modulation of non-invasive brain signals with force might be better captured non-linearly as in the invasive case (Naufel et al., 2019; Fatemi and Daliri, 2020). Hence, we interpret our non-linear approach to search for non-linear modulation signatures of non-invasive brain signals with graded bimanual forces.

6.2 Methods

6.2.1 Participants and task

Nine participants ($N = 9$) were asked to perform a bimanual isometric contraction task. The task used in this experiment was designed to study muscle-like and high-level characteristics of force control and their non-invasive decoding (Chapter 5). To do so, we provided the force profile with two characteristics (Fig. 6.1A rightmost). First, both hands were either contracted or relaxed at the same time. This allows us to distinguish between a relaxed and contracted state representing a high-level or coarse force-control state. Second, when the hands were in the contraction state, they had to track five force profiles (Fig. 6.1A left to right).

The first profile consisted of a constant contraction with both hands and was not used in this study. The remaining profiles (2 to 5, Fig. 6.1A) required each hand to track a different dynamic force target. Since each hand tracked a different target simultaneously, this part of the task represents muscle-like characteristics of force generation. These muscle-like characteristics are specifically associated with the group of muscles controlling each hand.

In summary, the task's high-level and muscle-like characteristics allow us to differentiate between relaxed and contracted bimanual states. Additionally, since the dynamic part was different for each hand, the task introduces the challenge of decoding simultaneous and hand-specific hand-grip generation. In particular, the different way each hand is engaged in dynamic force tracking gives a better representation of continuous control of force. Altogether, our task represents a controlled way of studying force generation during bimanual interactions.

The participants received visual feedback on the desired contraction trajectory that they were instructed to follow with each hand. Except for condition 1, in which both hands performed the same force profile, the remaining 2 to 5 conditions required contraction trajectories that were specific to each hand.

The trials in all conditions lasted 10 s, and all participants did 30 trials per condition, except two participants who were tired to perform the last block. The order of the conditions was randomised, but each condition was performed in blocks of 30 trials. For conditions 2 to 5, the highest level of contraction was 25% MVC, the lowest 10% MVC, and contraction average for all trajectories was 17.5% MVC. The trial's start was indicated with a "beep" sound, and for the first 2 s of conditions

2 to 5, the target indicated a flat 17.5% MVC that the subjects aimed to reach before the dynamic part of the task started. Each trial was followed by a randomised resting period uniformly distributed between 15 and 21 seconds, to avoid phasic constructive interference of systemic artefacts, e.g. Mayer waves and breathing, in the brain responses. The refreshing of the feedback on the screen was set to 100 Hz.

The task and the data used for the decoding experiments come from the Hybrid (fNIRS & EEG) bimanual isometric Grip dataset (HYGRiiP) dataset (Chapter 5). All participants were right-handed and healthy at the time of the study. Handedness was confirmed by the Edinburgh inventory (Oldfield, 1971) for all participants. The Imperial College Research Ethics Committee approved all procedures, and all participants gave their written informed consent. The experiment complied with the Declaration of Helsinki for human experimentation and national and applicable international data protection rules.

6.2.2 Recording setup and preprocessig

Twenty four (N = 24) co-aligned EEG and fNIRS channels covered the bilateral sensorimotor cortex and were used as the brain signals from which to decode the force signals generated with each hand (see Chapter 5 for details of the recordings and the setup).

The fNIRS signals were recorded using a NIRScout system (NIRx Medizintechnik GmbH, Berlin, Germany). We used a total of 12 channels per hemisphere (10 sources and 8 detectors in total) sampling at 12.5 Hz.

EEG was recorded using 24 channels of an ActiChamp amplifier (BrainProducts, Berlin, Germany) operating at 4 kHz (running software BrainVision, v1.20.0801). EEG was first downsampled to 500 Hz (with anti-aliasing down-pass filtering). Notch filters were applied at the mains (50 Hz) and fNIRS (12.5 Hz) frequencies and their harmonics. EEG was finally high-pass filtered above 1 Hz using a 5th order Butterworth filter. ICA was then applied to remove EOG artefacts rejecting 1 component when they had a correlation above 0.3 in absolute value.

Two grip force transducers (PowerLab 4/25T, ADInstruments, Castle Hill, Australia) were used to record the force generated by each hand, simultaneously recording at 1 kHz. The force signals were first resampled to 250 Hz, then band-pass filtered between 0.01 and 10 Hz with a Butterworth filter of order 3 and then again high-pass filtered with an elliptical filter of order 1 above 0.01 Hz. The

linear drift was removed per trial, and measures were finally converted to contraction values using the recorded MVCs before the experiment started.

The recorded data underwent additional processing before the decoding. Optical intensities were low-pass filtered below 0.25 Hz with a 7th order elliptical filter. Changes in haemoglobin concentration, were obtained using Eq. 3.2.

The processing of EEG was separated into two different streams, one generating 'raw' EEG measures and the other extracting time-courses of the power and phases of different EEG bands or, as we refer to them, 'Hilbert features' (Fig. 6.1C). These two streams allowed us to compare our algorithms' decoding capabilities when using 'raw' EEG data or its spectral content per frequency band ('Hilbert features'). When 'raw' EEG was used in the decoding, the remaining processing consisted of resampling all signals to 80 Hz. Otherwise, the EEG was decomposed in the power and phases of the following frequency bands as in Paek et al. (2019): delta (1 – 4 Hz), theta (4 – 8 Hz), alpha (8 – 13 Hz), beta (13 – 30 Hz), low-gamma (30 – 50 Hz), mid-gamma (70 – 110 Hz), high-gamma (130 – 200) using the Hilbert envelope method (Oppenheim, 1999) to generate the 'Hilbert features'. In this case, all signals and the newly computed EEG 'Hilbert features' were resampled to 12.5 Hz.

We also recorded the scalp skin's hemodynamic activity on the forehead using a NONIN 8000R (Tilburg, The Netherlands) in 17 of the sessions. The skin hemodynamics reflect the variations of oxygenated haemoglobin due to the systemic activity (pulse, breathing and Mayer waves) but do not contain brain hemodynamic responses (Gagnon et al., 2012). We used scalp hemodynamics to confirm that systemic variations in blood oxygenation at the skin level were not predictive of force.

After preprocessing all signals, epochs were extracted from 4 s before the "Go" instruction to 14 s after.

6.2.3 Decoders

We used six different decoders to compare their ability to decode bimanual force from non-invasive brain signals. In some cases, the decoders' main difference was the kind of EEG (raw or Hilbert features) used as input (Fig. 6.1C). Remaining signals, fNIRS and force, always underwent the same processing independently of the EEG. The only exception was the frequency at which we resampled the signals so all measures could be aligned in time. All non-raw models used signals resampled at 12.5 Hz, while raw models used signals resampled at 80 Hz.

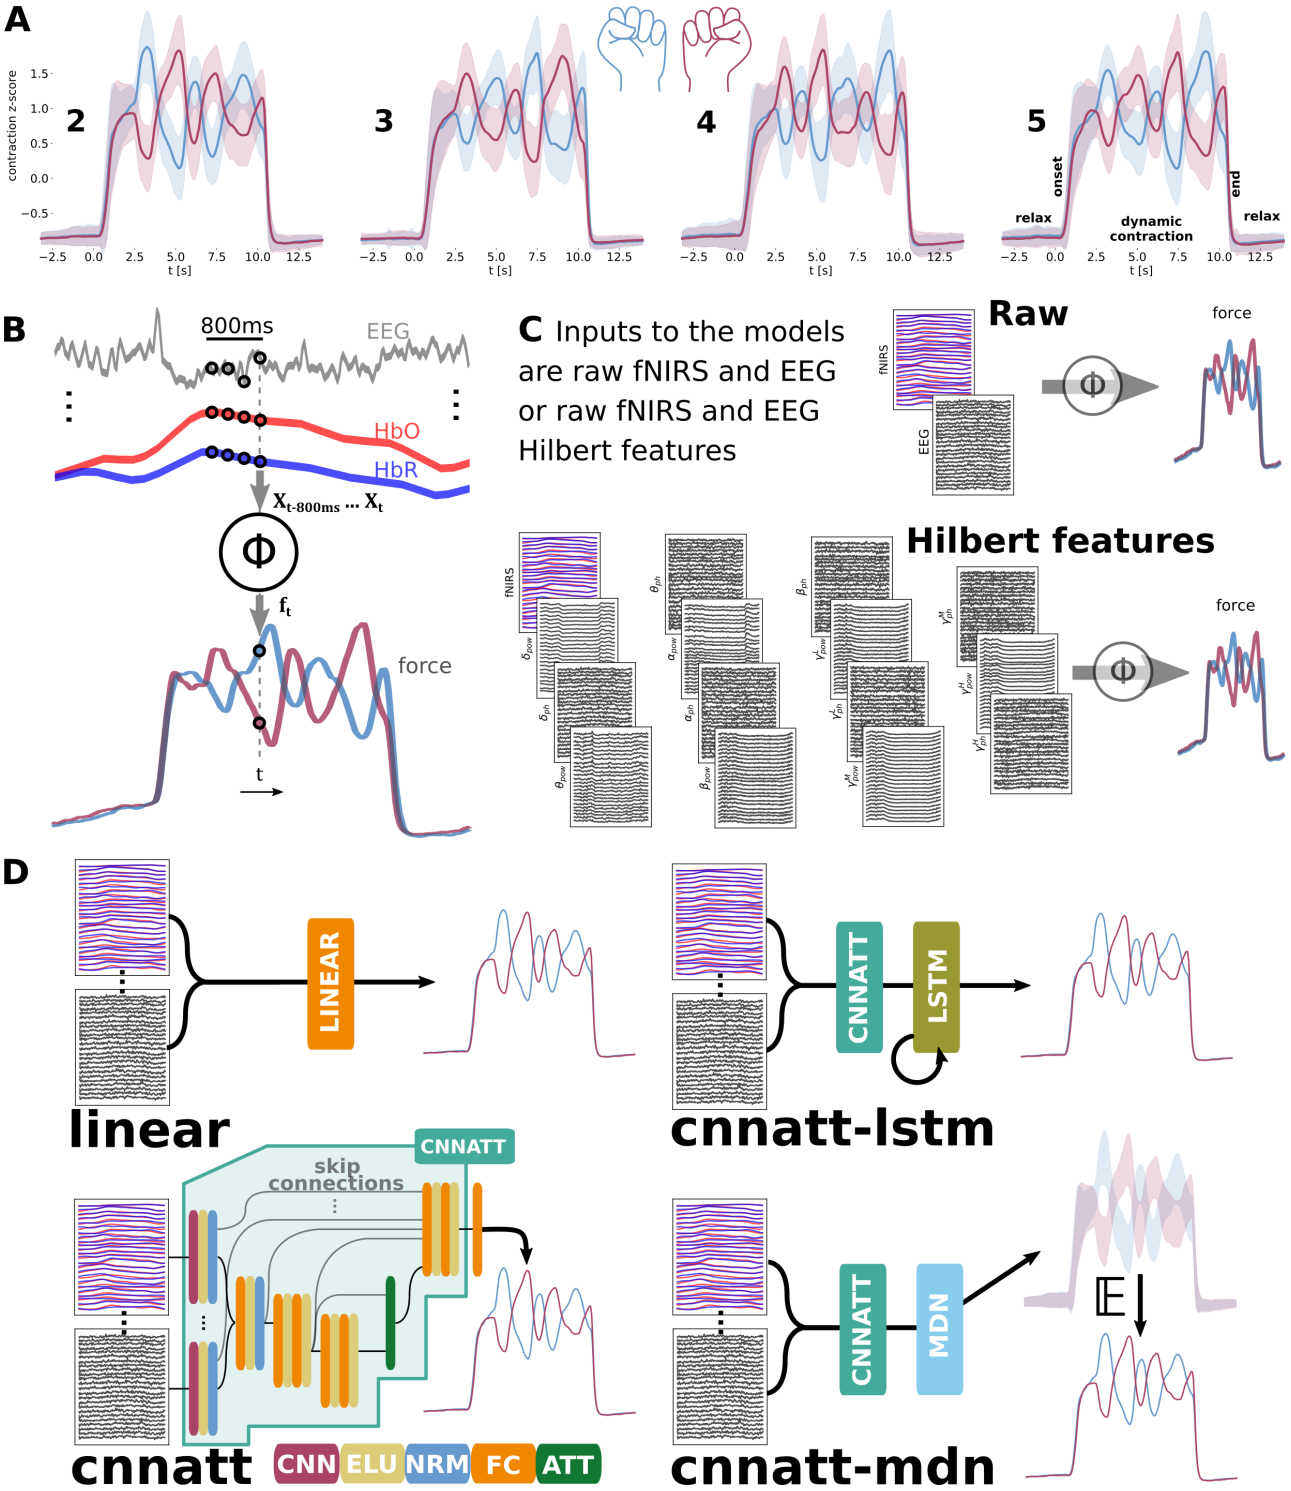


Figure 6.1: **A** Averaged and standard deviation of the contraction values (z-scored) across subjects four the force profiles used in this study. Both hands contracted ('onset') and relaxed ('end') simultaneously. **B** Decoding scheme. **C** Inputs. **D** Models used for force decoding. Modules are colour-coded: convolutional (CNN), exponential linear unit (ELU), layer-normalisation (NRM), fully connected (FC) and self-attention (ATT). Skip connections represent the activations that continue deep into the architecture and are also forked to the layers closest to the outputs. The *cnnatt* module ends in 20 hidden units in all cases, which then are transformed to 2 outputs using an FC, LSTM or MDN layer. *Cnnatt* and *cnnatt-mdn* were trained and tested with raw and Hilbert features inputs. The remaining models only used Hilbert features inputs.

We built these models with increasing complexity. From less to more complex the used models (Fig. 6.1D) were: (1) A *linear* model, (2) a convolutional neural network with residual modules and self-attention using 'Hilbert EEG features' (*cnnatt*) or 'raw' EEG (*cnnatt-raw*); (3) the last architecture using an LSTM layer in the output and only tested with 'Hilbert envelope features' (*cnnatt-lstm*) since EEG raw features delivered a high sampling frequency that considerably slowed down their training times and reduced the performance; (4) *cnnatt* with a Mixture Density Network (MDN) as output layer with 'raw EEG' or 'Hilbert EEG features' (*cnnatt-mdn-raw* and *cnnat-mdn*, respectively).

Every decoding scheme used 800 ms of brain signals history to decode the bimanual force (Figure 6.1B). The exception was *cnnatt-lstm* since, in theory, the recurrent architecture would allow for longer memories to be exploited in the decoding. Although most BCI decoders exploiting electrical signals use memory depths below 500 ms (Paek et al., 2019; Hayashi et al., 2019) we increased the memory to 800 ms. This way, we aimed to better capture the slower fNIRS signals without excessively increasing the input dimensionality. All inputs were z-scored using the channel-wise average and standard deviation values computed using the training data.

The decoding (Fig. 6.1B) can be expressed as

$$\hat{\boldsymbol{f}}_t = \phi(\boldsymbol{X}_{t-800\text{ms}}, ..., \boldsymbol{X}_t) \tag{6.1}$$

where $\hat{\boldsymbol{f}}_t$ represents the vector of bimanual left (L) and right (R) reconstructed forces $\hat{\boldsymbol{f}}_t = \left[\hat{f}_{\text{L},t}, \hat{f}_{\text{R},t}\right]^{\text{T}}$ and $\boldsymbol{X}_t$ the matrix of fNIRS and EEG features (either 'raw' or 'Hilbert envelope') at time t.

Previous invasive BCI research suggests that exploiting the task's temporal structure can be beneficial in force decoding (Sussillo et al., 2012; Naufel et al., 2019). We introduce recurrences in the model using an LSTM layer designed to capture temporal force dependencies. In particular, for the LSTM architecture (*cnnatt-lstm*) the memory-depth can be learned in a more flexible way, emphasised in the expression via the parameter τ. At the same time, LSTM allows the modelling to depend on previous decoded outputs $\hat{\boldsymbol{f}}$ via recurrence,

$$\hat{\boldsymbol{f}}_t = \phi(\boldsymbol{X}_{t-\tau}, ..., \boldsymbol{X}_t, \hat{\boldsymbol{f}}_{t-\tau}, ..., \hat{\boldsymbol{f}}_{t-1}). \tag{6.2}$$

Finally, the MDN layer was added to *cnnatt* (*cnnatt-mdn*) since invasive BCI evidence has reported improvements by including the variance and covariance structure of the task in the modelling (Sussillo

et al., 2012; Naufel et al., 2019). The task's covariance becomes particularly relevant in our decoding since we aim to decode two hands simultaneously. Indeed, hand activities are correlated during the relaxation, onset and end of force generation, and non-correlated or anti-correlated during dynamic force generation. MDN layers allow the force output to be modelled as a Gaussian mixture. In particular, we adopted the approach of Graves (2013) and explicitly modelled the bimanual force's covariance structure through a mixture of bivariate Gaussians using a maximum-likelihood loss during training. Hence, the MDN models (*cnnatt-mdn* and *cnnatt-mdn-raw*) represent the conditional probability of the force generated by the two hands given the 800 ms history of brain signals. The MDN were implemented as the weighted sum of bivariate Gaussians,

$$p(\boldsymbol{f}_t|\boldsymbol{X}_{t-800\text{ms}},...,\boldsymbol{X}_t) = \sum_{j=1}^{M} \pi_t^j \mathcal{N}(\boldsymbol{f}_t|\boldsymbol{\mu}_t^j, \boldsymbol{\sigma}_t^j, \rho_t^j), \tag{6.3}$$

with each bivariate Gaussian $\mathcal{N}$ modelled as,

$$\mathcal{N}(\boldsymbol{f}_t|\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t, \rho_t) = \frac{1}{2\pi\sigma_{t,\text{L}}\sigma_{t,\text{R}}\sqrt{1-\rho_t^2}} \exp\left[\frac{-Z_t}{2(1-\rho_t^2)}\right], \tag{6.4}$$

and

$$Z_t = \frac{(f_{t,\text{L}}-\mu_{t,\text{L}})^2}{\sigma_{t,\text{L}}^2} + \frac{(f_{t,\text{R}}-\mu_{t,\text{R}})^2}{\sigma_{t,\text{R}}^2} + \frac{\rho(f_{t,\text{L}}-\mu_{t,\text{L}})(f_{t,\text{R}}-\mu_{t,\text{R}})}{\sigma_{t,\text{L}}\sigma_{t,\text{R}}}. \tag{6.5}$$

Here, L and R indicate the left and the right hand emphasising the bivariate nature of our bimanual task. Specifically, two bivariate Gaussians ($M = 2$) were used with the aim to capture correlated and anti-correlated structures in the force outputs. Finally, the force decoded by *cnnatt-mdn* and *cnnatt-mdn-raw* can be computed as the expected value of the decoded force distribution,

$$\hat{\boldsymbol{f}}_t = \mathbb{E}\left[p(\boldsymbol{f}_t|\boldsymbol{X}_{t-800\text{ms}},...,\boldsymbol{X}_t)\right]. \tag{6.6}$$

All models followed the same training and testing strategy. In contrast to hemCNN, where the models were tested across subjects (Chapter 4), here we train and test models per subject since the relationship between the brain signal and the hand contraction of subjects with different strengths might not share a similar mapping.

In order to introduce enough variance in our decoding task, we train and test in all task conditions presenting dynamic force generation (2 – 5 in HYGRiiP). For each subject, we follow a 5-fold training

approach. Once a fold was left-out for testing, we used a 90% of the remaining data to train the neural networks and a 10% for validation, while the *linear* model used 100% of the data for training. We ensured that the representation of each task condition was balanced across data splits. The validation data was used to stop the training when the validation did not continue decreasing.

The models using MDN used a maximum likelihood approach for training. The rest used the Mean Squared Error (MSE) loss between the predicted and the actual force. Once the models finished training, they were tested, and we gathered these test results across testing folds for statistical comparison.

6.3 Results

6.3.1 Comparison of EEG, fNIRS and multimodal (EEG & fNIRS) performance in force trajectory reconstruction

We first want to know whether the use of a multimodal (fNIRS & EEG) system brings any benefit in terms of bimanual isometric force trajectory decoding. To do this, we use the *linear* method in Paek et al. (2019). They used this *linear* method to decode continuous unimanual force from EEG. Here, we expand their approach to fNIRS and the multimodal system to decode our bimanual force task.

We quantify the decoding performance using two metrics: the Spearman's Correlation Coefficient (SCC) (Zwillinger and Kokoska, 1999) (Fig. 6.2B), and Fraction of Variance Accounted For (FVAF) (Fig. 6.2A) between the reconstructed force trajectories and the corresponding hand real ones. In particular, FVAF is a more stringent measure than correlation and has been used before to evaluate force trajectory reconstruction performance (Fagg et al., 2009; Flint et al., 2016; Naufel et al., 2019). FVAF has values between $(-\inf, 1]$ with 1 representing a 100% reconstruction and 0 representing a 0% reconstruction, which is as bad as using the average of the actual trajectory as the prediction. Negative FVAF values are even worse than using the average of the signal for reconstruction. FVAF can be computed as a percentage using,

$$\text{FVAF} = 100 \cdot \left(1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2}\right)\% \tag{6.7}$$

with y_i the i^{th} sample of the real signal, $\hat{y}_i$ the corresponding sample of the predicted signal, $\bar{y}$ the average of the real signal and N the number of samples in the signal. From Eq. 6.7 we can observe

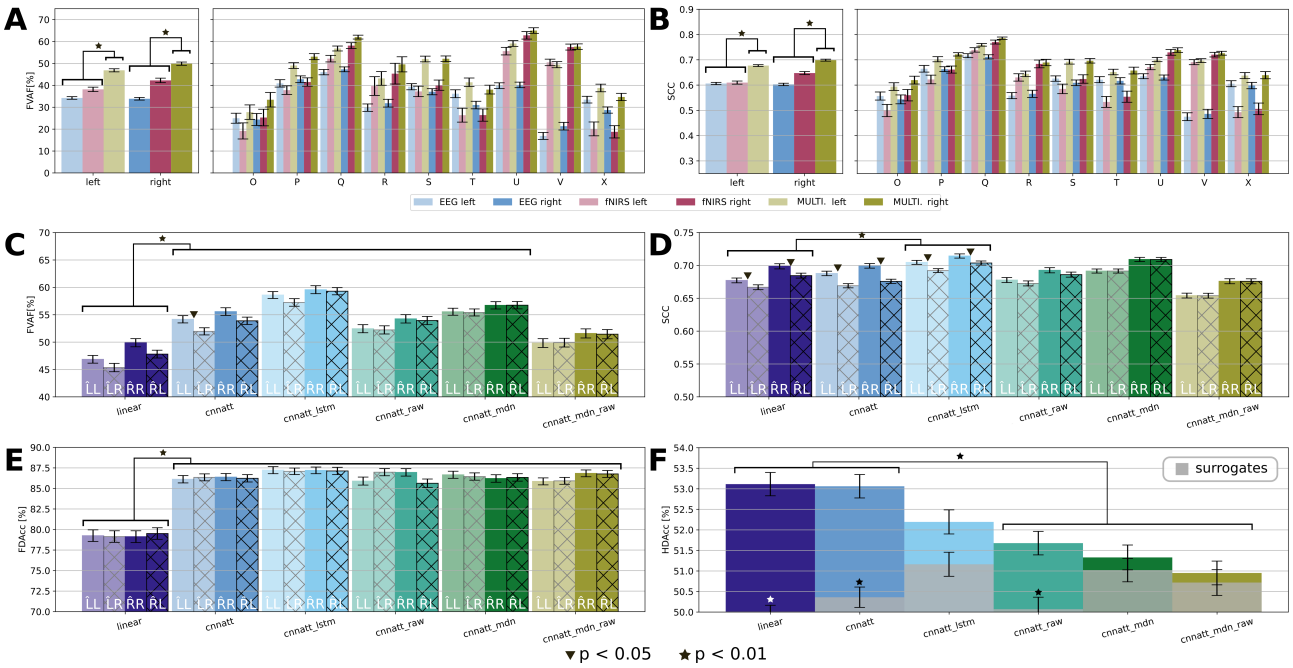


Figure 6.2: Force reconstruction performance metrics. Error bars represent the standard error of the mean (1020 trials; Tukey's honest significance test ($\blacktriangledown p < 0.05$ and $\bigstar p < 0.01$). Performances computed between the predicted hand and the corresponding real hand (plain bars, ĹL and R̂R) or the opposite real hand (shaded bars, ĹR and R̂L). **A** Spearman's correlation coefficient (SCC) of the *linear* baseline model for EEG, fNIRS and multimodal measures. (Left) The multimodal *linear* system achieves higher SCC than EEG or fNIRS alone (see Table 6.4). **B** Ibidem for the fraction of variance accounted for (FVAF[%], Table 6.5). **C** SCC for each method. The right hand is overall better decoded than the left (see Tables 6.2 and 6.6). **D** FVAF for each method. Deep-fusion methods generally outperform the *linear* method (see Tables 6.2 and 6.6). **E** Force detection accuracy (FDAcc) per model. Deep-fusion methods outperforms the *linear* method. **F** Hand Disentanglement Accuracy (HDAcc) per model and their surrogates in grey (see Table 6.2).

that FVAF also has a more direct interpretation in terms of decoding error: Namely, we can interpret the quotient as an MSE value that is normalised to the actual trajectory variance. In contrast to MSE, because the quotient is subtracted to 1, the highest the FVAF the better the decoding.

The SCC is defined as the Pearson correlation coefficient between the ranks extracted from the raw measures $\hat{y}_i$ and y_i, respectively $r\hat{y}_i$ and ry_i,

$$\text{SCC} = \frac{\text{cov}(r\hat{y}_i, ry_i)}{\sigma_{r\hat{y}_i}\sigma_{ry_i}} \tag{6.8}$$

With cov(), the covariance and σ the variance of the ranks.

The results of this analysis are shown in Figures 6.2A and 6.2B and Tables 6.5 and 6.4. At the

	EEG	fNIRS	Multim.	Max. Increase
SCC	0.597/0.594	0.606/0.642	0.676/0.698	0.1
FVAF	34/33	38/42	47/50	17

Table 6.1: Cross-subject average of force trajectory reconstruction measures, left/right-hand SCC and FVAF, for the linear model using EEG, fNIRS or both (multimodal). At the population level and for both hands and measures, the multimodal approach achieves significant higher performance than either signal (EEG or fNIRS) alone (see Table 6.5).

population level, the increase of reconstruction performance of fNIRS (SCC = 0.610[L]/0.648[R]; FVAF[%] = 38.2[L]/42.3[R]%) and EEG (SCC = 0.606[L]/0.603[R]; FVAF[%] = 34.2[L]/33.9[R]%) is significantly lower than that of the multimodal system (SCC = 0.677[L]/0.699[R]; FVAF[%] = 46.9[L]/49.9[R]%; Tukey's honest significance test, α = 0.01, Fig. 6.2A). In a per-subject basis, the multimodal system significantly outperforms both the EEG and the fNIRS signals in 3/18 cases for SCC (Table 6.4) and 6/18 for FVAF[%] (Table 6.5).

Finally, using a similar *linear* model, we predicted the force exclusively using the skin's haemodynamics during force generation. The force predicted with scalp hemodynamics was significantly less predictive of the real force (SCC = −0.07 and FVAF[%] = −1%) than the force predicted with fNIRS (SCC = 0.5; FVAF[%] = 32.5%; t-test, pval < $1e-6$). In particular, force reconstructed from skin haemodynamics show negligible force reconstruction performances.

6.3.2 Comparison of *linear* and deep-fusion methods for bimanual force trajectories reconstruction

Linear models have shown a modest continuous force reconstruction capability from EEG (Paek et al., 2019) and fNIRS (Nambu et al., 2009). While invasive studies have shown the advantage of adding non-linearities in force decoding (Naufel et al., 2019) approaches like DL remain to be explored in continuous force decoding from non-invasive signals. In particular, these approaches might offer additional advantages in fusing (deep-fusion) EEG and fNIRS with different physical and physiological origins (Ball et al., 1999).

To evaluate the decoding capabilities of DL in bimanual force continuous decoding, we focus on three aspects of force generation. We first evaluate the continuous force reconstruction capabilities (force trajectory reconstruction). Then, we proceed to identify the aspects of this continuous reconstruction that are better decoded. In particular, we analyse how well models can distinguish a high-level force generation characteristic. Namely, how well they can distinguish between a relaxed and a contracted state (force detection accuracy). Then, we introduce a hand decoding specific measure that measures the models' ability to detect what hand is being more contracted at each time step (hand disentanglement accuracy, HDAcc).

Force trajectory reconstruction, SCC and FVAF. We now interrogate the capabilities of DL to reconstruct bimanual force trajectories, particularly how automatic feature extraction methods compare to using linear approaches.

We show that DL models outperform linear models in terms of FVAF (Fig. 6.2C). In contrast, SCC is less able to convey these differences (Fig. 6.2D) and most of them are non-significant (Tukey's honest significance test, $\alpha = 0.05$; Table 6.2). Additionally, when the same models extract features from the raw signals, the performance decreases significantly compared to using the 'Hilbert features' (i.e. *cnnatt* and *cnnatt-mdn* respectively compared to *cnnatt-raw* and *cnnatt-mdn-raw*; Tukey's honest significance test, $\alpha = 0.05$; Table 6.2). In terms of SCC and FVAF, the right hand is also significantly better decoded than the left hand by all models (t-test, $\alpha = 0.05$). The models' reconstruction performance is in the $0.65 - 0.72$ range for SCC and $47 - 60\%$ for FVAF.

Furthermore, in terms of SCC, all models preserve a significant sensitivity to the hand except for those using raw measures and the MDN output layer. In terms of FVAF, the only model achieving this is the *cnnatt* for the left hand. More importantly, the best performers in terms of SCC and FVAF, are *cnnatt* and *cnnatt-lstm* due to the higher SCC, FVAF, Force Decoding Accuracy (FDAcc) and comparable Hand Disentanglement Accuracy (HDAcc) to the *linear* model (Tables 6.6 and 6.2). However *cnnatt-lstm* seems to saturate to flat values of contractions for the dynamic part of the force generation (Fig. 6.3, column 3).

Although SCC and FVAF are valuable measures to rank decoding performance and test for hand sensitivity, it is hard to picture what is achievable in terms of decoding. Figure 6.3 presents representative examples for all subjects (rows) and methods (columns). On the other hand, the dynamic part of the bimanual contraction is only reconstructed in some examples and to a limited extent. While the reconstruction somewhat captures some of the bimanual dynamics, it struggles to reach the required amplitudes. We further analyse these aspects in the following sections.

Force Detection Accuracy, FDAcc. Force Decoding Accuracy (FDAcc) measures how well the models detect that a hand is producing force or is relaxed. For each time step, we consider that the hands are producing force if the signal is above the threshold separating the two amplitude populations of forces, low and high levels. The low level represents relaxation (values close to zero). We apply Equation 4.4 to compute the accuracy, this time using the label 0 for low force level and the label 1 for high force level. We also consider the hand-specificity of this measure. We do this by computing

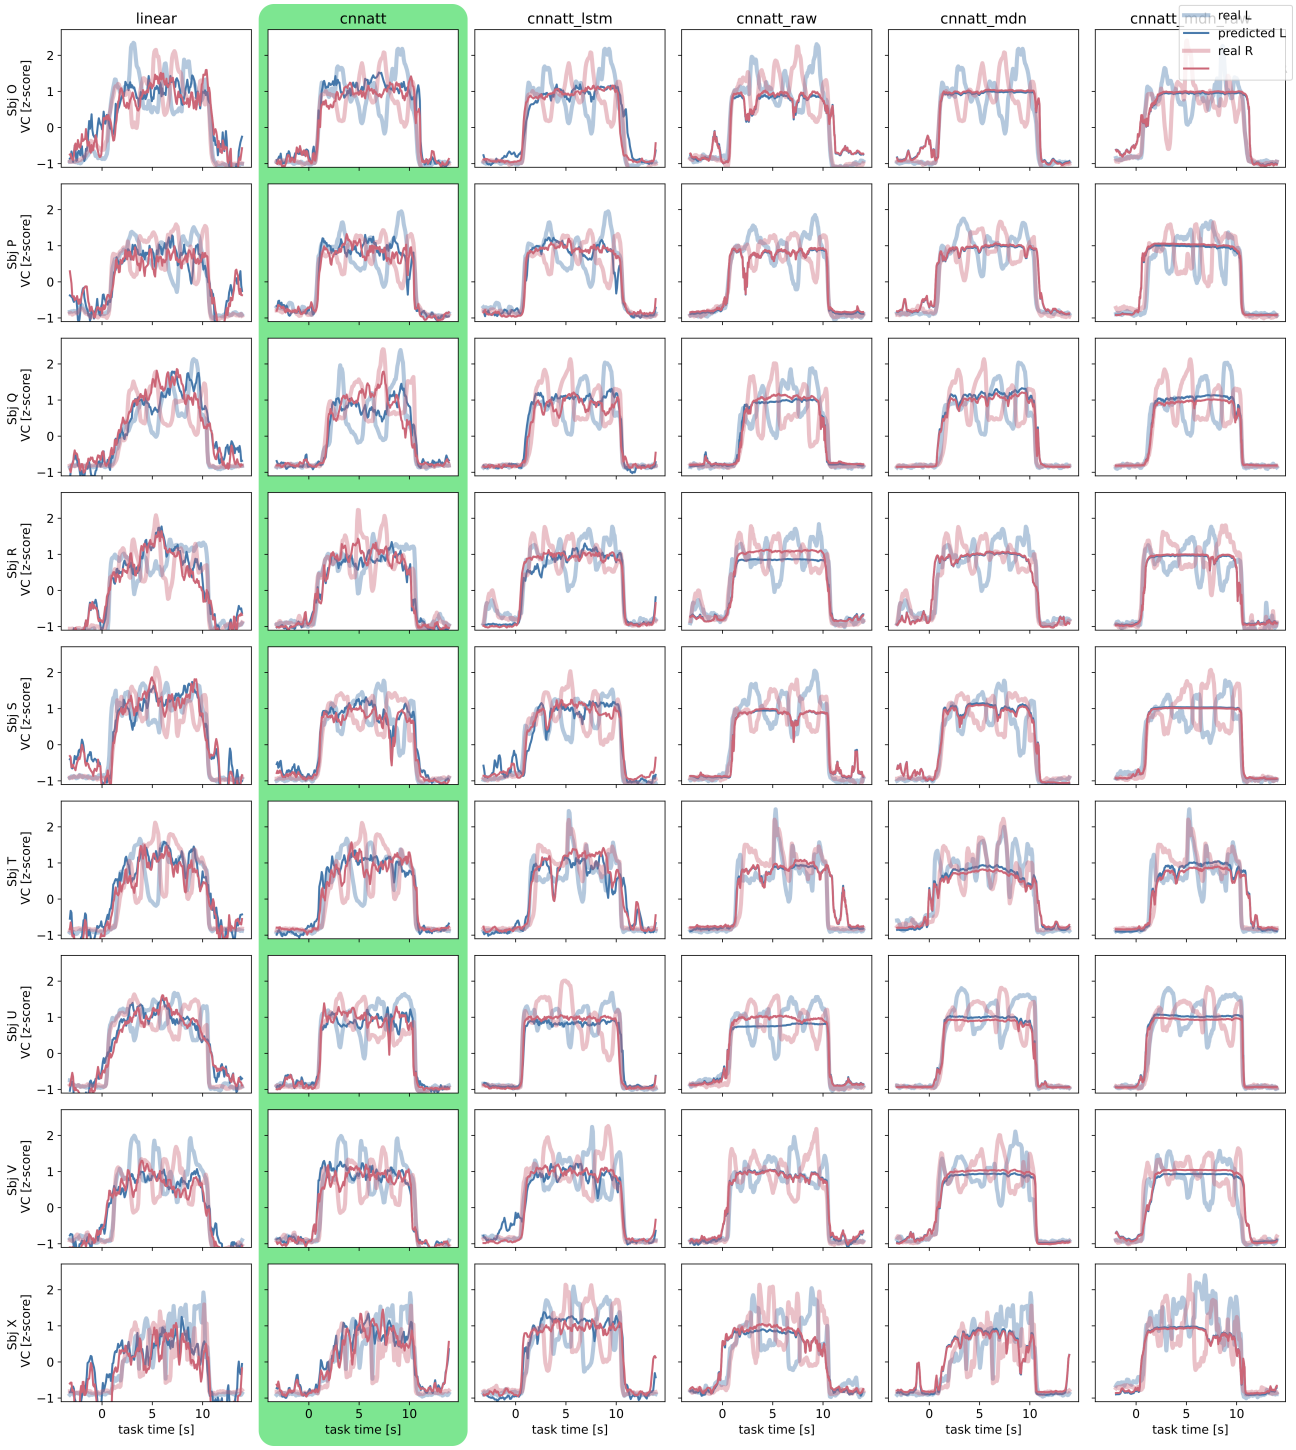


Figure 6.3: Representative examples of decoded force trajectories per subjects in columns. The feature most accurately decoded corresponds to the beginning and end of the contraction for all methods. As reflected by the HDAcc values the dynamic changes of forces during force production are captured by the *linear* and the *cnnatt* methods only to a limited extent (see Fig. 6.2F and Tables 6.6 and 6.2). In some cases, the profile is more or less well decoded, but the hand it belongs to is confused. In other cases, the hand is accurately captured, but the amplitude does not reach the required value. *Cnnatt-raw*, *cnnatt-mdn* and *cnnatt-mdn-raw* notably struggle at disentangling the hands as captured by the HDAcc measure in Figure 6.2F. We select *cnnatt* (green column) as the best model due to its balanced performance between force decoding and hand disentanglement (see main text).

		cnnatt		cnnatt lstm		cnnatt raw		cnnatt mdn		cnnatt mdn raw	
		L	R	L	R	L	R	L	R	L	R
linear	SCC	0.255	0.900	0.001	0.025	0.900	0.829	0.043	0.310	0.001	0.001
	FVAF[%]	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.033	0.576
	FDAcc	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	HDAcc	0.900		0.219		0.006		0.001		0.001	
cnnatt	SCC			0.008	0.037	0.315	0.751	0.900	0.390	0.001	0.001
	FVAF[%]			0.001	0.002	0.507	0.785	0.706	0.879	0.001	0.002
	FDAcc			0.664	0.010	0.900	0.900	0.900	0.900	0.900	0.846
	HDAcc			0.279		0.010		0.001		0.001	
cnnatt lstm	SCC					0.001	0.001	0.078	0.900	0.001	0.001
	FVAF[%]					0.001	0.001	0.025	0.072	0.001	0.001
	FDAcc					0.636	0.030	0.900	0.014	0.384	0.001
	HDAcc					0.782		0.284		0.030	
cnnatt raw	SCC							0.060	0.016	0.001	0.013
	FVAF[%]							0.022	0.176	0.086	0.114
	FDAcc							0.900	0.900	0.900	0.658
	HDAcc							0.900		0.486	
cnnatt mdn	SCC									0.001	0.001
	FVAF[%]									0.001	0.001
	FDAcc									0.765	0.801
	HDAcc									0.900	

Table 6.2: Summary of bimanual force decoding methods. Green indicates significant better performance of the column method compared to the row method. Red indicates significant worse performance. Grey indicates non-significant differences (Tukey's honest significance test, $\alpha = 0.05$).

FDAcc between each hand's predicted trajectory and the actual trajectory of the hand the model was not trained to decode (Fig. 6.2E).

DL models outperform the *linear* model in detecting when the hands are producing force (Tukey's honest significance test, $\alpha = 0.05$; Tables 6.6 and 6.2). However, when we use the opposite decoded hand to detect the force of the actual hand, the levels of accuracy (Fig. 6.2E shaded bars) are similar to those using the corresponding hand (plain bars in Fig. 6.2E; Table 6.6). We anticipated this since we designed our task to reflect differences in decoding due to different dynamic force generation per hand. In particular, we asked participants to simultaneously contract both hands, which explains why there are no hand-specific differences in the detection of force. From a force detection point of view, we consider *cnnatt-lstm* as the best method given its higher FDAcc (Tables 6.6 and 6.3).

Hand Disentanglement Accuracy, HDAcc. We showed that it is possible to reconstruct around a FVAF $\approx 55\%$ (*cnnatt*) of the force trajectories from multimodal signals (EEG & fNIRS). While the detection of the onset and end of the force is very accurate (FDAcc > 85%), we observed that the dynamic part of the force was less well reconstructed, particularly in terms of amplitude (Fig. 6.3).

To have a better idea of the models' ability to capture hand specific information during the dynamic part, we use the Hand Disentanglement Accuracy (HDAcc). In particular, HDAcc measures how accurately the models can decide which hand is producing higher force than the other at each time step. We apply again Equation 4.4 with labels 0 for the left hand applying higher force and 1 for the left hand applying higher force. Figure 6.2F shows the average HDAcc achieved by each method.

The *linear* and the *cnnatt* models have significantly higher HDAcc compared to the rest of the models except *cnnatt-lstm* (Tukey's honest significance test, $\alpha = 0.05$; Tables 6.6 and 6.2). Nonetheless, *cnnatt* has a higher HDAcc than *cnnatt-lstm*. We empirically computed the random level of HDAcc by shuffling the fNIRS and EEG inputs in the temporal dimension ("surrogates" in Fig. 6.2F). The empirical random level was situated within the $50 - 51.5\%$ HDAcc range depending on the model, and only the *linear*, *cnnatt* and *cnnatt-raw* models had significantly higher HDAcc than the random empirical level corresponding to their models (Tukey's honest significance test, $\alpha = 0.01$).

Comparison summary. A summary of the comparison is presented in Table 6.2 for all the methods and all the scoring measures (Tukey's honest significance test, $\alpha = 0.05$). The table shows how *cnnatt* and *cnnatt-lstm* significantly outperform the linear model in at least two scoring measures but perform similarly in terms of HDAcc. Nonetheless, *cnn-lstm* does not have an HDAcc significantly higher than that of its surrogate, and it is lower as that of *cnnatt*.

The lower HDAcc of *cnnatt-lstm* is also qualitatively captured in the corresponding column of Fig. 6.3 which shows a relatively flat decoding of the dynamic forces. In addition, *cnnatt* is the only method that preserves hand-specificity for both hands in terms of SCC and the right hand in terms of FVAF (Table 6.6).

Altogether, we consider *cnnatt* as the most appropriate method to decode bimanual force trajectories from non-invasive multimodal (fNIRS and EEG) measures while *cnnatt-lstm* can offer better detection of forces with less hand-specificity.

Model interpretation The interpretation of the models used to decode force can provide insight into how the brain controls the generation of force and what features from the non-invasive signals provide the highest information. We take a double-sided approach. We first perform a sensitivity analysis to see what features are more important for force decoding in a causal sense. Next, we reveal the brain activity patterns that the selected model (*cnnatt*) associates with each hand during force decoding, also known as task encoding (Weichwald et al., 2015).

Figure 6.4 shows the sensitivity results per feature, hand and performance measure. Our perturbation approach consists of randomly shuffling an input feature in the time dimension and then measuring its effect on the decoding. We use the same performance measures as before. In particular, this test evaluates how important the features' causal time-course is in decoding forces. To standardise

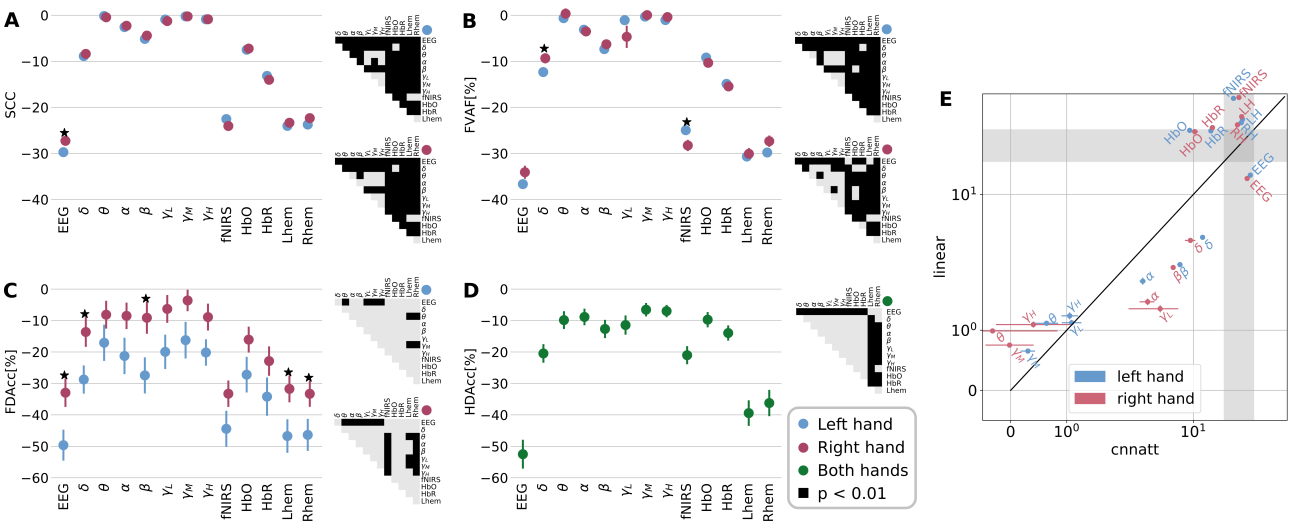


Figure 6.4: **A B C D** *Cnnatt* sensitivity analysis. A 0% corresponds to the performance when inputs are not perturbed, and a −100% change corresponds to the performance when all inputs are perturbed. Coloured circles represent mean values and bars standard error of the mean. We use Tukey's honest significance test to compare means. In **A** and **B**, EEG is more important than fNIRS. The delta band is more important than any remaining band (t-test, $*\alpha < 0.01$). **E** *Linear* model and *cnnatt* sensitivities for each input feature. For visualisation, values are multiplied by −1 and shown in a logarithmic scale. The horizontal and vertical grey bars represent the same range of sensitivity on both axes. *Cnnatt* shows a more balanced performance in EEG and fNIRS (within the vertical grey bar range) than the *linear* model (outside the horizontal grey bar range, 20 – 90%). *Cnnatt* is also more sensitive than the *linear* model to the EEG bands with more saliency, delta, beta and alpha.

the measures, we compute the per cent change in performance with a 0% change representing the performance of the unperturbed signals and −100% the maximum performance reduction when all features are perturbed. As we can see, the perturbation of any of the features leads to a reduction in performance (negative changes) in all measures, providing evidence that the input's causal order is essential in the decoding.

Trajectory decoding performance metrics, SCC and FVAF, provide similar results to each other (Fig. 6.4A and B). The perturbation of all EEG features leads to the most significant reduction of performance for both measures, except for the right hand for FVAF. For the latter, the change in performance between EEG and fNIRS is not significant (Tukey's honest significance test, $\alpha = 0.01$).

The second most significant reduction in FVAF and SCC is due to the perturbation of fNIRS. Within the EEG bands, the perturbation of the delta band leads to higher performance reduction, followed by the beta and then the alpha bands. The remaining bands, theta and all gamma bands have minimal effect and do not seem to play a significant role in force decoding. In terms of fNIRS, HbR has a more definitive role in the decoding of force compared to HbO, and the perturbation of fNIRS leads to a lower performance decrease in decoding compared to EEG. However, this decrease is more notable when compared to each of the EEG bands. The perturbation of all the signals belonging to the left or

right hemisphere leads to a similar decoding performance decrease, with no hemisphere seeming to be more important than the other in providing information for the decoding of force.

The effect of the perturbation of features on the force detection accuracy (FDAcc, Fig. 6.5C) losses most of the differences found for FVAF and SCC. However, it reveals differences across hands that, to some extent, could already be observed in FVAF or SCC. Namely, in terms of FDAcc, the perturbation of EEG and, in particular, the delta and beta bands significantly worsen the left-hand's decoding more than that of the right-hand (t-test, α = 0.01). The same can be observed when the signals of each hemisphere are perturbed separately. Regarding force detection, this indicates that the information available in the left hand is somewhat contingent on the information available for the right hand. Overall, the detection of right-hand contractions dominates in the brain signals regardless of the hemisphere and the signal type.

In terms of hand decoding accuracy (HDAcc, Fig. 6.5D), the perturbation of all the EEG bands together lead to the most significant reduction of accuracy followed by the perturbation of the left or the right hemisphere (Tukey's honest significance test, α = 0.01). However, there are no significant differences between the different bands of EEG and fNIRS.

The previous analysis focuses on measuring *cnnatt's* sensitivity to the perturbation of its inputs' temporal order. It provides evidence that the decoding is causal and sheds some light on the essential inputs for force decoding. However, it does not reveal much how the brain encodes force generation. To reveal the temporal and spatial brain dynamics that accompany bimanual force generation, we analyse the saliency of the inputs (Simonyan et al., 2013). The inputs' saliency provides a measure of how important each part of the input is in *cnnatt* force decoding. Specifically, for a given loss, the saliency specifies the direction in the input space in which a change would lead to a greater loss decrease. More precisely, the saliency is the gradient of the model parameters with respect to the loss evaluated at the inputs values. We compute the saliency for each hand separately in two different ways. The first computes each hand's saliency when the evaluated hand produces more force than the other during dynamic force generation (Figs. 6.5A, B and Fig. 6.7). The second way is related to the saliency associated with each hand individually, whatever the force the hands generate (Fig. 6.5C and Fig. 6.6).

In Figures 6.5A and B, we can observe the time course of every feature's saliency during the task. The peaks are indicated in the figure, and we can observe similar results as in the inputs perturbation test.

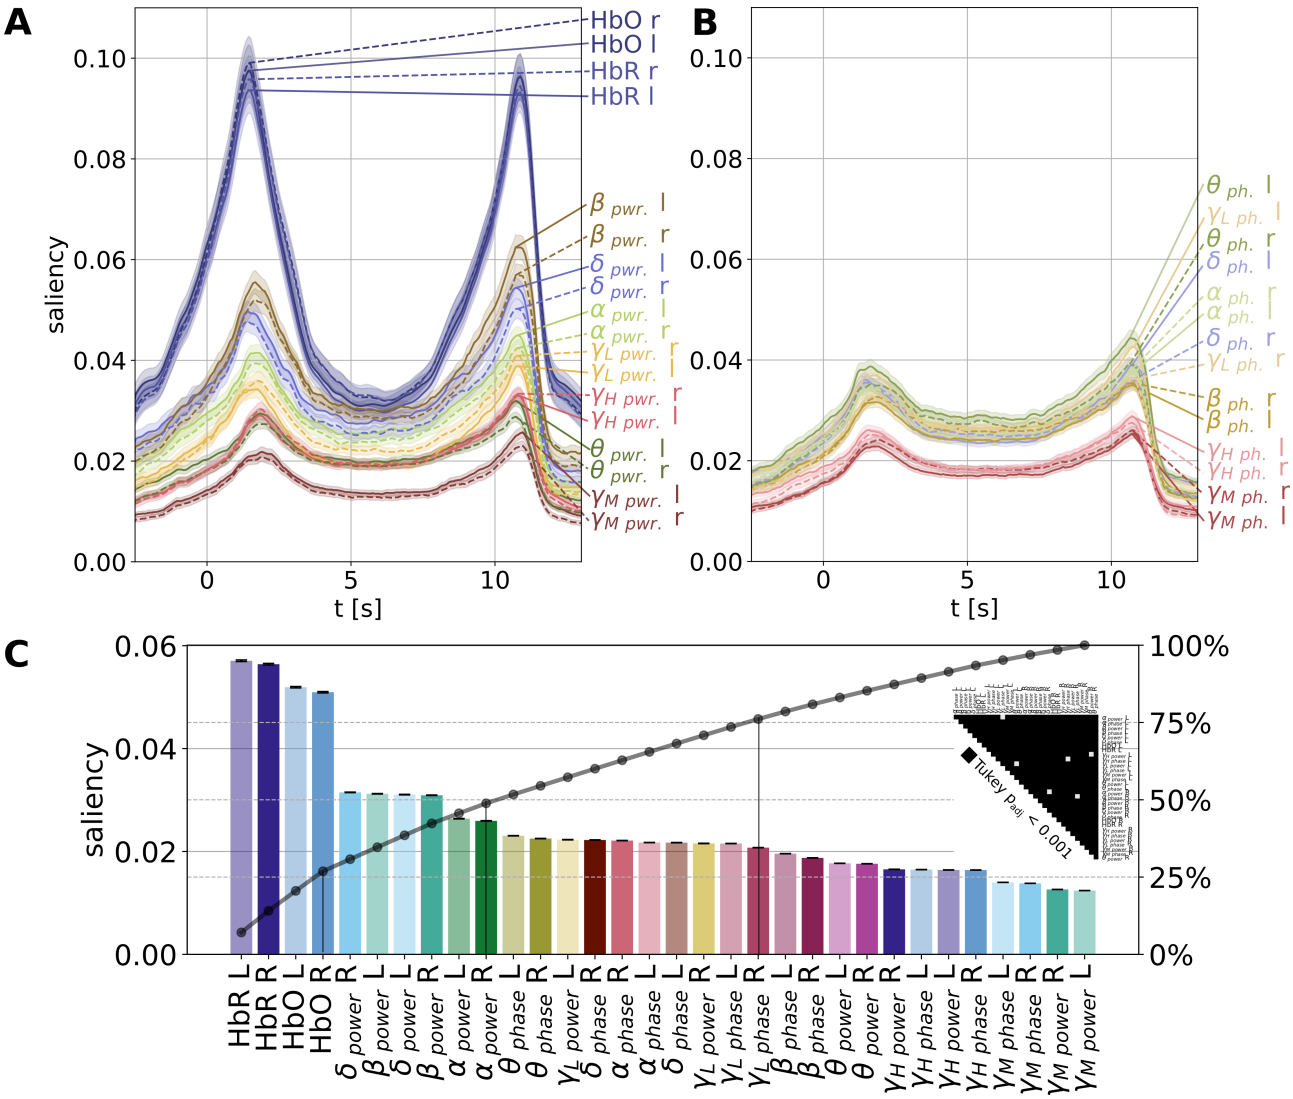


Figure 6.5: Saliency time course. **A** Mean and standard error of the saliency for each hand associated with each input modality. **B** Ibidem for the phases. The saliency has clear peaks at the beginning and end of the dynamic generation of force in the power and phase features. **C** Mean and standard error associated with each hand per input modality during the dynamic part of force generation in decreased order (left y-axis). In the right axis, the percentage of accumulated saliency is represented. fNIRS corresponds to 25% of the total saliency. A 50% is reached when delta, beta and alpha power are included. A 75% is reached when the phase of theta, alpha and delta, and the power of low-gamma are included. Note that fNIRS has higher saliency than any of the EEG bands. However, the saliency of all EEG bands represents a 75% of the total, with the rest corresponding to fNIRS.

Namely, for fNIRS and EEG power features, we can observe that the saliency of HbO and HbR is higher than that of any of the individual EEG features (6.5A). Additionally, among EEG features, we can observe beta's higher saliency, then delta and then alpha. These are finally followed by low-gamma, high-gamma and finally theta power.

Interestingly, this order changes for the phases of the features, with theta phase being more salient than that of beta, inverting the order of their power saliency (Fig. 6.5B). This saliency order is almost intact when the saliency is computed during the dynamic part, using the time-steps where the hand under analysis produces higher force than the other (Tukey's honest significance test, $\alpha = 0.001$; Fig. 6.5C). This figure also shows the percentage of total saliency accounted for by each feature. By itself, fNIRS accounts for 25% of the saliency, which is more than any of the individual EEG Hilbert features. However, when we put together all EEG features, they account for the remaining 75% of saliency.

The higher saliency associated with EEG than with fNIRS is similar to what the *cnnatt* perturbation analysis conveyed. Mainly, that EEG is more important in the decoding when all bands are considered, but fNIRS alone is more important than each EEG band separately. Interestingly, this is different to what we observed in the *linear* model which delivered better decoding accuracy for fNIRS than EEG (Figs. 6.2A 6.2B).

In contrast to what happens during the whole task, in the dynamic part, HbR has higher saliency than HbO. Finally, Figure 6.5C highlights the order of relevance of the features during the dynamic part. As we said, 25% of the saliency corresponds to fNIRS. A 50% would include delta, beta and alpha power, and a 75% would also add the phases of theta, delta and alpha, but not beta, and the power of low-gamma. Following this order of feature saliency, we will next analyse the temporal and spatial patterns of saliency.

Figure 6.6 shows the topographical distribution of saliency associated with the left and the right-hand when each hand produces a higher force than the other. The heat-map represents the saliency value while the blue dots represent channels with significantly higher saliency for one hand than the other (t-test, $\alpha = 0.001$).

Remarkably, despite the low accuracy in HDAcc, the use of each hand leads to disentangled activity patterns at the cortical level. Some of the patterns are purely contralateral, others are ipsilateral, and others show specific bilateral activations. HbR shows a specific significant contralateral saliency for the right hand while HbO shows it for the left and somewhat more spread. Delta power shows ipsilateral

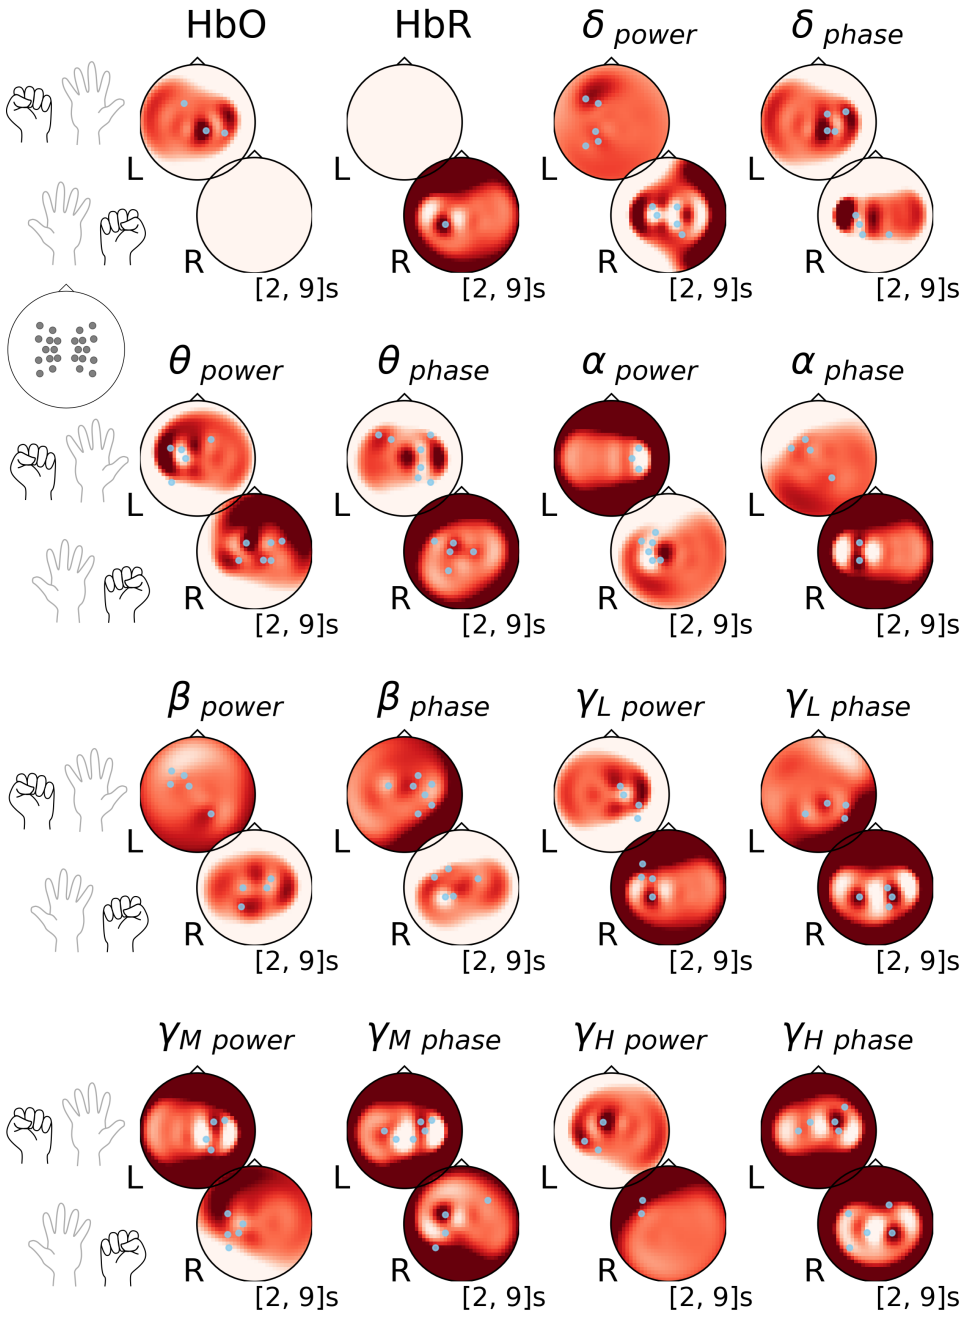


Figure 6.6: Topoplot of saliency associated with each hand during the dynamic part of force generation. The saliency is computed for each hand, and a t-test is performed to compare the level of saliency across hands for each channel. The channels with significant higher saliency for one hand than the other are represented in blue, and their saliency included in the topoplot (t-test, $p < 0.001$).

activity for the left and bilateral for the right hand, with delta phases showing stronger contralateralisation than the former. Beta power shows a bilateral pattern with ipsilateral frontal projections and contralateral dorsal projections of saliency. Finally, alpha power shows pure contralateralisation, similar to low-gamma.

Some of these patterns survive when the saliency is computed for each hand regardless of the force applied. Figure 6.7 shows the saliency time-course for each feature and associated with each hand during the task. If we focus on the dynamic part of the task, $[2, 10]$ s, we can see how some of the patterns identified for each hand in the dynamic part (Fig. 6.6) survive when the saliency is for each hand regardless of the force it produces (Fig. 6.7). Namely, the contralateralisation of alpha power is preserved, delta shows a similar bilateral pattern, and beta maintains a bilateral pattern with ipsilateral frontal projection.

We further inspect the representation of force in non-invasive brain signals in Figure 6.8. We compute the lateralisation index (LI) of the saliency per pair of symmetric channels around C3 and C4. The saliency is computed during aligned dynamic changes of force (left-most column, Fig. 6.8) across conditions and with opposite force trajectories per hand. The figure shows that some channel pairs and features have a significant modulation that follows the bimanual force generated to a limited extent. For example, for channel 9, the HbO saliency has a similar modulation as the force generated with an initial right hemisphere dominance for the left hand (blue), generating more force at the beginning. This right hemisphere dominance is then followed by a higher left hemisphere dominance for the left hand when the left-hand force is lower than the right-hand force. We can observe similar patterns for other features and channels, especially for beta power.

6.3.3 *Linear* vs *cnnatt* comparison.

Finally, we compare the *linear* and the selected DL model, *cnnatt*. In particular, we want to understand how each method exploits the information in the inputs to decode force trajectories. We follow the same previous time-perturbation analysis and compare the FVAF[%] performance decrease for the *linear* and *cnnatt* models. For readability, we multiply the decrease in performance values by -1 to get positive values. For a given model, these values represent how much the model's decoding depends on a given feature. Specifically, a 0% indicates that the model does not depend on that feature, while a 100% indicates it depends on that feature to decode the force. The results are shown on the left in

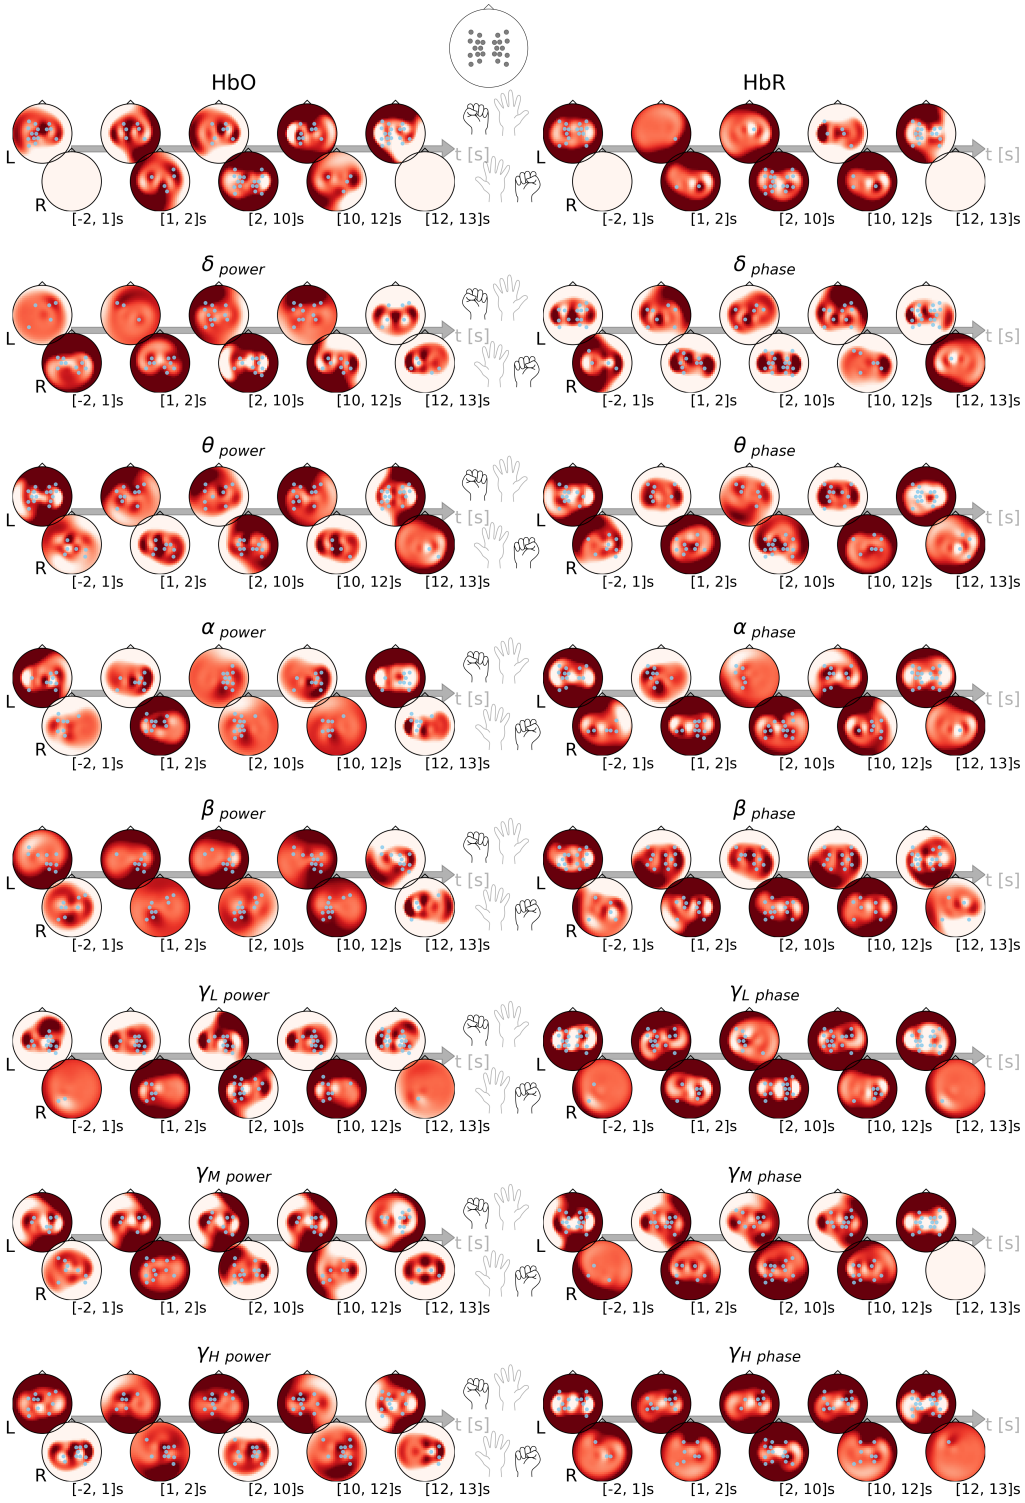


Figure 6.7: Saliency topoplot associated with each generated force. The saliency is computed using the error associated with the decoded force, which is more stringent since both forces are used to generate each hand saliency, and differences across hands emerge only from the different dynamic hand engagement in the task. In contrast to Fig. 6.6, the saliency here is computed during the whole trial in specified chunks of time. Some patterns observed in Fig. 6.6, in particular, delta, alpha and beta, can also be observed here.

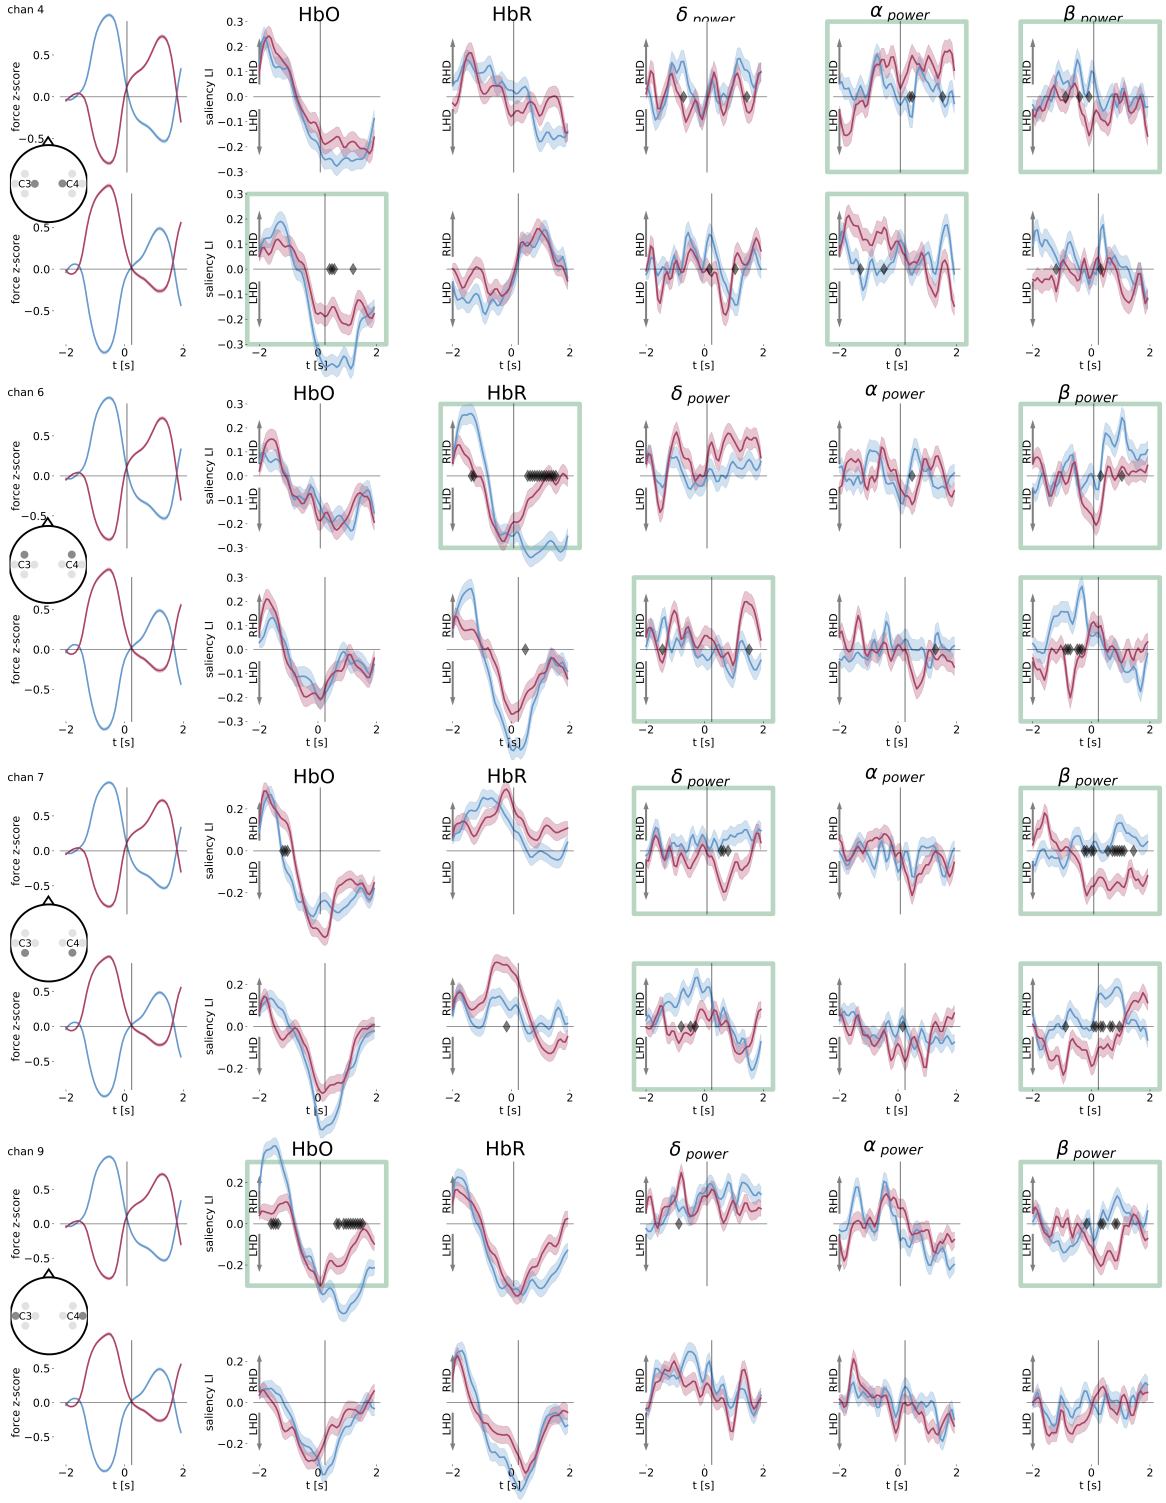


Figure 6.8: Saliency average time course during the dynamic generation of bimanual force. The leftmost column shows the average and standard deviation of the z-score contraction with peaks and valleys aligned. We use the saliency time-course associated with each shown input (HbO, HbR and delta, alpha and beta power) to compute the lateralisation. Except for the leftmost column, in the remaining columns, positive y-axis values indicate right hemisphere dominance (RHD, upwards arrow) and vice versa for the left hemisphere dominance (LHD, downwards arrow). Diamonds represent significant differences between the lateralisation of the saliency associated with each hand (t-test, $p < 0.05$). Blue represents the left hand and red the right hand. In green selected examples.

Figure 6.4E.

As we can see, *cnnatt* depends more in EEG and the relevant EEG bands than the *linear* model, while the *linear* model depends more in fNIRS (Fig. 6.4E). The vertical and horizontal grey bars in the figure are used as visual guides marking the same interval of dependence in both axes ($20-90\%$). For *cnnatt*, the perturbation of EEG or fNIRS leads to a decrease in force reconstruction within the range indicated by the vertical grey bar. The opposite can be observed for the *linear* model, where perturbations in EEG lead to lower reconstruction decrease than for fNIRS (both outside the horizontal grey bar interval).

Indeed, the *linear* model presented a worse decoding with EEG (SCC = $0.606[L]/0.603[R]$; FVAF[%] = $34.2[L]/33.9[R]\%$) than fNIRS (SCC = $0.610[L]/0.648[R]$; FVAF[%] = $38.2[L]/42.3[R]\%$, Tukey's honest significance test, $\alpha = 0.01$, Fig. 6.2A). This indicates that the *cnnatt* strikes a more balanced dependence between EEG and fNIRS than the *linear* model which mainly depends on fNIRS.

Note that the EEG in this case consists of the Hilbert envelopes in all EEG bands, and, therefore, it includes signals changing at different frequencies. Since the *linear* model performs worse with EEG and the *cnnatt* model not only depends more on EEG but also has a more balanced performance between EEG and the slower fNIRS, this suggests the *cnnatt* model can improve the fusion of signals with different time-scales.

6.4 Discussion

In this chapter, we set out to explore EEG and fNIRS deep-fusion to reconstruct bimanual force trajectories on a time-step basis. To ensure that we exploit brain activities, we rejected ocular sources of activity using ICA and confirmed the negligible decoding power of skin haemodynamic activity.

We first found an increase of 0.1 SCC and 17% FVAF[%] when combining fNIRS and EEG (Table 6.1), showing that the *linear* fusion of multimodal brain signals outperforms each single modality at the population level. For every subject, the multimodal system led in performance which was significant in 17% (SCC) and 33% (FVAF[%]) of the cases, corresponding to subjects with balanced EEG-fNIRS performance (Tables 6.4 and 6.1). In the remaining cases, the multimodal system performed as the best subsystem but doubled the chances of acquiring valuable signals.

Our results expand previous multimodal work, showing that combining EEG and fNIRS improves the

classification of force trials (Hansen et al., 2014; Fu et al., 2016; Li et al., 2017). In particular, our multimodal *linear* force reconstruction (0.70-0.79 SCC, 30-65 FVAF[%]) improved previous reconstructions using EEG (0.42-0.51 SCC) (Paek et al., 2019) and was similar to that of a non-causal fNIRS decoder (32-68 FVAF[%]) (Nambu et al., 2009).

The reconstruction (FVAF[%]: *cnnatt* 54.9%, *linear* 48.5%) and force detection (FDAcc: *cnnatt* 86.3%, *linear* 78.9%) were further improved by our *cnnatt* method, performing similarly in correlation (SCC *cnnatt* 0.694, *linear* 0.688) and hand disentanglement (*cnnatt* and *linear* 53.1% HDAcc, Tables 6.2, 6.6). *Cnnatt-lstm* reached higher reconstruction (60% FVAF[%]) and force detection (87% FDAcc), but led to lower hand disentanglement (52.3% HDAcc, non-significantly different from that of its surrogates, Fig. 6.2F, Table 6.2 , 6.6).

Overall, our EEG-fNIRS deep-fusion methods outperformed the multimodal *linear* fusion. Nonetheless, the hand disentanglement accuracy was much lower (53.1% HDAcc, 50-51% empirical baseline) than the detection of force (*cnnatt-lstm* 87% FDAcc, *linear* 78% FDAcc, Fig. 6.2E). Additionally, most saliency corresponded to the initiation and completion of forces (Fig. 6.5A and B), suggesting that most of the reconstruction gain was due to the ability to detect and model sharp changes of force, as it has been invasively observed in grasping (Pistohl et al., 2008, 2013). Our deep-fusion approach (*cnnatt-lstm*, 87% FDAcc) favourably compares to a previous attempt at detecting force in a time-step basis that achieved a 74.8% accuracy using EEG and a *linear* decoder (Paek et al., 2019).

Interestingly, although by task design the force detection is not hand-specific (see Table 6.6), when the delta or beta bands, or the left or right hemispheres, were perturbed, force detection decreased more for the left than for the right hand (t-test, $\alpha = 0.01$; Fig. 6.4C). The left hand's higher sensitivity suggests that the force generation detection is hand-specific. Additionally, since the left hand was more sensitive to both hemispheres perturbation, the right-hand seems to be more widespread across hemispheres than the left. However, while many real-world manipulations require both hands contracting simultaneously (Lima et al., 2020; Haar and Faisal, 2020), this aspect is rarely tested, and most studies focus on unilateral tasks (Li et al., 2017; Trakoolwilaiwan et al., 2017; Chiarelli et al., 2018; Khan and Hasan, 2020; Ghonchi et al., 2020). In contrast, our EEG-fNIRS deep-fusion method offers more robust real-time grasp decoding, reflecting in real-world tasks which need decoders immune to hand cross-talk.

Several reasons explain the better performance of our deep-fusion approach. First, the *linear* model performs worse with EEG than with fNIRS (Table 6.1), reflecting on its poor capacity to combine

the different EEG frequency bands and its greater dependence on the slowest fNIRS signals (Fig. 6.4E). In contrast, *cnnatt* is more dependent on EEG (Fig. 6.4A and B, 75% of total saliency, Fig. 6.5C) and strikes a more balanced cross-modality dependence than the *linear* model (Fig. 6.4E). The different physiological origins of EEG and fNIRS lead to very different signal velocities whose coupling require extensive feature engineering (Li et al., 2017; Khan and Hasan, 2020). Our deep-fusion method provides better decoding by automatically finding the best temporal features across modalities.

Second, to further exploit these time-scale differences, we added an LSTM layer to *cnnatt*. The addition of the LSTM layer translated into better force detection but worse hand disentanglement. Additionally, explicitly modelling the bimanual covariance (*cnnatt-mdn*) led to a reduction of hand disentanglement (Table 6.2, Fig. 6.2F). We think that these architectures' internal mechanisms might be biasing the models to detect the force edges. These edges are shared across hands and might be driving the gradient through all time-steps during backpropagation through time in LSTM or in the cross-hand covariance modelling in MDN, masking more subtle hand-specific features.

Third, compared to previous DL classification approaches using raw fNIRS (Trakoolwilaiwan et al., 2017; Ortega and Faisal, 2021), EEG (Schirrmeister et al., 2017; Ortega et al., 2018; Gemein et al., 2020) or both (Chiarelli et al., 2018; Ghonchi et al., 2020), we early realised the benefit of including residual (He et al., 2016) and attention modules (Vaswani et al., 2017) to reconstruct force in a time-step basis. Without them, we could not outperform the *linear* method in any metric.

Finally, using raw EEG reduced the hand disentanglement performance compared to using Hilbert features (Table 6.2). Raw EEG required a sampling frequency that increased 20 times the temporal input dimensionality. Since the output is redundantly over-sampled, this raw EEG approach may require approximately 20 times more examples to start learning temporal features that outperform the Hilbert transform.

Remarkably, where previous studies were limited to unimanual (Paek et al., 2019) or a left/right sequential (Nambu et al., 2009) tasks. Despite the limited hand-specificity effect in the decoding (53% HDAcc), our deep-fusion method can dissociate hand-specific cortical activities during simultaneous hand use.

First, we find a contralateral dominant activation for the left hand in HbO and a similar contralateral activation for the right hand in HbR ($[2, 9]$ s in Fig. 6.7). These cortical patterns suggest that the network learns hand-related contralateral activations across haemoglobin types, similar to those previously found during unimanual tasks (Shibuya et al., 2016; Ortega and Faisal, 2021).

In the EEG delta band, there is a clear frontoparietal ipsilateral pattern for the left hand that becomes more bilateral for the right. Since the delta band is the most critical EEG band in our decoding (Figure 6.4E), a dominant cortical representation of the dominant hand might explain why the dominant hand is generally better decoded than the non-dominant left hand (Table 6.1, 47/50% L/R FVAF[%] and 0.676/0.698 L/R SCC). This kind of low frequency frontoparietal bilateral activations have been previously found in unilateral grasping (Di Bono et al., 2015; Iturrate et al., 2018; Paek et al., 2019). In particular, the ipsilateral component is more prevalent during fine motor control (Iturrate et al., 2018), reflecting the activity of a network in charge of modulating efferent motor commands with sensory afferent information (Olivier et al., 2007; Castiello and Begliomini, 2008; Mulliken et al., 2008; Andersen and Cui, 2009; Davare et al., 2011). Our bimanual task shows that this kind of patterns survives during bimanual control, supporting evidence of bilateral hemispheric interactions during bimanual control (Serrien et al., 2006) and a dominant cortical representation of the dominant hand (Volkmann et al., 1998).

Alpha shows a clear contralateralisation (Fig. 6.6 and 6.7), suggesting that during bimanual control forces mostly rely in the contralateral hemisphere as in the unimanual case (Penfield and Boldrey, 1937; Pfurtscheller et al., 2006).

Beta reveals interesting fronto-parietal patterns across hemispheres (Fig. 6.6 and 6.7). Namely, there is a frontal projection in the ipsilateral hemisphere and a dorsal projection in the contralateral. While previous unimanual studies revealed bilateral frontoparietal patterns as separated entities for the dominant hand (Iturrate et al., 2018; Paek et al., 2019), our bimanual task suggests that in beta, these fine control networks span across hemispheres. Activity in alpha and beta has been previously associated with top-down control (Klimesch et al., 2007; Fries, 2015; Sadaghiani and Kleinschmidt, 2016). Our results suggest that beta may also serve a hand coordination function across hemispheres (Serrien et al., 2006; Ames and Churchland, 2019).

Note that our deep-fusion method mostly relies in the aforementioned delta, beta and alpha bands and fNIRS for the decoding (Figure 6.4A and B), assigning similar amounts of saliency to each modality (25% for delta, beta and alpha and 25% for HbO and HbR, Figure 6.5C).

Despite lower saliencies in the gamma band, contralateral gamma activities (Ball et al., 2008) are also revealed (Fig. 6.6) but do not survive when saliency is computed regardless of the produced force (Fig. 6.7). In general, higher frequencies have lower saliencies (Fig. 6.5C) in line with previous invasive

research showing that high-frequency bands are more attenuated in EEG (Buzsáki et al., 2012) and more relevant for sensorimotor feedback (Flint et al., 2016; Branco et al., 2019a,b).

Interestingly, the theta band was the only low-frequency band with more saliency in its phase than in its power (Fig. 6.6). Theta phase has been related to behaviour monitoring and errors (Wang et al., 2005; Womelsdorf et al., 2010) in the medial and lateral prefrontal cortex (Cavanagh et al., 2009). Here, the left-hand theta-phase saliency is more frontal than the right-hand saliency ([2, 10] s; Fig. 6.7), suggesting that the left hand is more subject to be monitored, resulting in higher frontal engagement.

Overall, our results suggest that some unimanual cortical activities are also present in the bimanual case and can be detected non-invasively with DL. However, differences in cortical resource allocation for each hand can reflect differences in hand decoding error.

Moreover, *cnnatt* showed hand-specific lateralised saliencies that followed the time-course of the force produced by the corresponding hand (Fig. 6.8). These modulation traces suggest that some of the minor hand disentanglement effects we observed (53% HDAcc, 50 – 51.5% random level) are captured by the non-invasive techniques and enhanced at the population level. Previous studies discarded the linear modulation of non-invasive brain signals with force (Paek et al., 2019; Hayashi et al., 2019). Moreover, invasive approaches provide evidence that neuronal activity is non-linearly modulated with graded forces (Ethier et al., 2016; Naufel et al., 2019). Our results suggest that these non-linear modulations could be detected non-invasively but are still impaired by the low SNR of non-invasive approaches (Ball et al., 2009; Buzsáki et al., 2012). Next-generation non-invasive sensors, combining photonics, electrical fields and acoustics (Torricelli et al., 2014; Phan et al., 2016; De Roever et al., 2017; Tang et al., 2016) could increase the SNR and benefit from the deep-fusion methods that we here advance to unlock real-time, continuously graded and bimanual control in non-invasive BMI.

6.5 Conclusion

We showed that multimodal approaches, together with DL regression techniques, are valuable tools in bimanual force decoding. Our results can be extended for BCI force detection decoding and provide the first step in bimanual force decoding from non-invasive brain signals. We could decode force generation periods with high accuracy and some force trajectories to a satisfactory extent. However,

hand-specific muscle-like characteristics of force control remain still challenging to decode in a robust manner. Nonetheless, we provide evidence suggesting that this kind of information reaches the scalp even if it is more affected by noise.

Here we demonstrated that a multimodal (EEG and fNIRS) system is better than either of them separately. We provide additional evidence that hand-specific motor control signals can reach the scalp in the form of optical and electrical signals. On the one hand, our multimodal results motivate the development of more complex sensing systems. On the other hand, our DL methods set a benchmark for multimodal fusion in the context of force trajectory reconstruction.

Altogether, multimodal sensing and DL decoding show a promising combination that can lead to significant breakthroughs in non-invasive BCI. Indeed, the next generation of non-invasive brain recording systems is becoming increasingly complex in the physical principles they utilise. These include the combination of different physical phenomena (e.g. photonics, acoustics and electrical fields), different physiological phenomena (e.g. respiratory cellular activity, haemodynamic responses) and cerebral information (i.e. morphology and activity). Since our methods are data-driven and agnostic physical measures used, they could be easily applied to emergent non-invasive brain function sensing modalities.

Our results support the idea that technical advances in machine learning and multimodal brain-sensing systems may soon materialise in more robust non-invasive BCI. As in our task, this may allow for muscle-like and 2-DoF continuous control. Such 2-DoF continuous control could favourably impact people who have lost bimanual or unimanual hand control. In either case, this control could enable the recovery of bimanual interactions for people with different degrees of disability. While invasive approaches have already demonstrated capabilities to do this (e.g. Pistohl et al., 2008), non-invasive BCI opens up new opportunities for people who might benefit from a temporary use of the BCI, for instance, during rehabilitation or ambulatory use (Casson et al., 2010), or do not have the means or are not eligible for costly and hazardous surgeries.

More directly, our results can have their most significant impact in the development of BCIs and robot-assisted rehabilitation of grip since we provide a robust and interpretable way of decoding hand-grip detection.

Supplementary tables

	precision [%]	sensitivity [%]	specificity [%]	F1 score
linear	95.5	86.4	87.1	0.907
cnnatt	96.1	91.4	88.4	0.937
cnnatt-lstm	97.0	92.8	90.9	0.949
cnnatt-raw	96.7	90.3	90.6	0.934
cnnatt-mdn	96.3	92.1	88.7	0.941
cnnatt-mdn-raw	97.4	90.1	92.4	0.936

Table 6.3: Statistical measures for the binary detection of force (relaxation vs contraction). All deep models outperform the linear method in force detection (light grey) and have the highest performance (green). In particular, the *cnnatt-lstm* presents the highest F1 score and is the best choice for force detection purposes.

Subject	Hand	EEG	fNIRS	Multim.	Tukey adjusted p value		
					EEG/fNIRS	EEG/Multim.	fNIRS/Multim.
O	Left	0.560	0.495	**0.597**	0.045	0.351	0.001
	Right	0.554	0.557	**0.628***	0.900	**0.014**	**0.019**
P	Left	0.664	0.621	**0.705**	0.073	0.088	**0.001**
	Right	0.660	0.658	**0.718***	0.900	**0.001**	**0.001**
Q	Left	0.709	0.737	**0.767**	**0.026**	**0.001**	0.094
	Right	0.698	0.770	**0.786**	**0.001**	**0.001**	0.254
R	Left	0.550	0.630	**0.647**	**0.001**	**0.001**	0.585
	Right	0.555	0.680	**0.689**	**0.001**	**0.001**	0.895
S	Left	0.606	0.580	**0.686***	0.440	**0.001**	**0.001**
	Right	0.589	0.621	**0.690**	**0.032**	0.101	0.069
T	Left	0.613	0.523	**0.648**	**0.001**	0.270	**0.001**
	Right	0.611	0.534	**0.652**	**0.005**	0.200	**0.001**
U	Left	0.619	0.667	**0.697**	**0.001**	**0.001**	0.068
	Right	0.612	0.729	**0.738**	**0.001**	**0.001**	0.784
V	Left	0.456	0.689	**0.697**	**0.001**	**0.001**	0.862
	Right	0.468	0.718	**0.726**	**0.001**	**0.001**	0.882
X	Left	0.601	0.488	**0.635**	**0.001**	0.312	**0.001**
	Right	0.602	0.498	**0.640**	**0.001**	0.280	**0.001**
POP	Left	0.597	0.606	**0.676***	0.369	**0.001**	**0.001**
	Right	0.594	0.642	**0.698***	**0.001**	**0.001**	**0.001**

Table 6.4: Comparison of SCC force reconstruction mean values (1080 measures) per hand and between EEG, fNIRS and multimodal (fNIRS & EEG) systems (Tukey's honest significance test, $\alpha = 0.05$). In bold, the highest SCC, which always correspond to the multimodal system. The asterisk (*) indicates where the multimodal system is significantly better than the EEG and fNIRS systems used separately (Tukey's honest significance test, $\alpha = 0.05$). At the population level (POP), the differences are significant (Tukey's honest significance test, $\alpha = 0.05$).

Subject	Hand	EEG	fNIRS	Multim.	Tukey adjusted p value		
					EEG/fNIRS	EEG/Multim.	fNIRS/Multim.
O	Left	26.1	19.5	**29.9***	0.262	0.632	**0.038**
	Right	26.3	25.3	**35.3**	0.900	0.107	0.064
P	Left	41.4	37.6	**49.8***	0.223	**0.001**	**0.001**
	Right	42.5	40.9	**52.9***	0.759	**0.001**	**0.001**
Q	Left	45.0	51.8	**56.7***	**0.001**	**0.001**	**0.023**
	Right	45.6	57.7	**61.8***	**0.001**	**0.001**	**0.035**
R	Left	29.4	39.7	**43.8**	**0.041**	**0.002**	0.584
	Right	30.9	45.0	**49.8**	0.091	**0.001**	0.570
S	Left	37.0	36.8	**51.3***	0.900	**0.001**	**0.001**
	Right	34.6	39.5	**51.2**	0.121	**0.001**	**0.001**
T	Left	35.6	25.4	**41.3**	**0.007**	0.214	**0.001**
	Right	30.4	24.8	**37.5**	0.168	0.058	**0.001**
U	Left	37.8	54.7	**58.7**	**0.001**	**0.001**	0.129
	Right	38.1	62.1	**64.6**	**0.001**	**0.001**	0.419
V	Left	16.4	49.8	**49.9**	**0.001**	**0.001**	0.900
	Right	20.2	56.8	**57.7**	**0.001**	**0.001**	0.880
X	Left	33.2	19.2	**38.6**	**0.001**	0.236	**0.001**
	Right	29.2	18.0	**34.7**	**0.001**	0.160	**0.001**
POP	Left	33.5	37.7	**47.1***	**0.001**	**0.001**	**0.001**
	Right	33.0	41.6	**49.8***	**0.001**	**0.001**	**0.001**

Table 6.5: Comparison of FVAF force reconstruction mean values (1080 measures) per hand and between EEG, fNIRS and multimodal systems (Tukey's honest significance test, $\alpha = 0.05$). In bold, the highest SCC average values. These highest SCC values always correspond to the multimodal system. The asterisk (*) indicates where the multimodal system is significantly better than both the EEG and the fNIRS system alone (Tukey's honest significance test, $\alpha = 0.05$). At the population (POP) level the differences are significant (Tukey's honest significance test, $\alpha = 0.05$).

Method	Case	FDAcc	p-adj	SCC	p-adj	FVAF	p-adj	HDAcc
linear	$\hat{\text{L}}$L/$\hat{\text{L}}$R	78.7/78.3	0.704	0.677/0.667*	0.046	46.9/45.4	0.148	53.1
	$\hat{\text{R}}$R/$\hat{\text{R}}$L	79.0/79.0	0.967	0.699/0.684*	0.009	49.9/47.8	0.105	
cnnatt	$\hat{\text{L}}$L/$\hat{\text{L}}$R	86.1/86.3	0.713	0.688/0.669*	<0.001	54.2/52.0*	0.018	53.1
	$\hat{\text{R}}$R/$\hat{\text{R}}$L	86.4/86.2	0.816	0.699/0.676*	<0.001	55.6/54.0	0.084	
cnnatt	$\hat{\text{L}}$L/$\hat{\text{L}}$R	87.0/86.9	0.864	0.705/0.692*	0.003	58.6/57.2	0.149	52.3
lstm	$\hat{\text{R}}$R/$\hat{\text{R}}$L	87.1/87.2	0.807	0.714/0.704*	0.017	59.6/59.3	0.778	
cnnatt	$\hat{\text{L}}$L/$\hat{\text{L}}$R	85.9/87.0	0.099	0.678/0.673	0.342	52.5/52.2	0.834	51.7
raw	$\hat{\text{R}}$R/$\hat{\text{R}}$L	87.0/85.6	0.054	0.693/0.686	0.226	54.3/54.0	0.773	
cnnatt	$\hat{\text{L}}$L/$\hat{\text{L}}$R	86.7/86.4	0.732	0.691/0.691	0.997	55.5/55.4	0.885	51.3
mdn	$\hat{\text{R}}$R/$\hat{\text{R}}$L	86.2/86.3	0.838	0.709/0.709	0.990	56.7/56.7	0.980	
cnnatt	$\hat{\text{L}}$L/$\hat{\text{L}}$R	85.9/85.9	0.912	0.654/0.654	0.974	49.8/49.9	0.946	51.0
mdn raw	$\hat{\text{R}}$R/$\hat{\text{R}}$L	86.8/86.8	0.905	0.676/0.676	0.989	51.6/51.5	0.907	

Table 6.6: Method comparison. The asterisk (*) indicates hand-specific significant differences (t-test, $\alpha = 0.05$). Metrics are computed between decoded and real left-hand ($\hat{\text{L}}$L), decoded left and real right-hand ($\hat{\text{L}}$R), decoded and real right-hand ($\hat{\text{R}}$R) and decoded right and real left-hand ($\hat{\text{R}}$L). We use the performance to reconstruct the opposite real hand to test for hand-specific decoding. *Cnnatt* is selected due to better FDAcc and SCC and FVAF and similar and significant HDAcc than the *linear* model.

Chapter 7

Conclusion

7.1 Thesis achievements in context

This thesis's work focuses on the development of Deep Learning (DL) force decoding algorithms for non-invasive Brain-Computer Interface (BCI). These interfaces are beneficial for users who have a partial or complete deterioration of their natural brain-muscle communication pathway (Wolpaw et al., 2000), helping them to assist (Onose et al., 2012), rehabilitate (Daly and Wolpaw, 2008; Soekadar et al., 2011) and even restore (Bouton et al., 2016; Ganzer et al., 2020) their ability to interact with the environment. To put this research in the context of clinical needs, we focused on the generation of hand forces which is the most missed function among people with Spinal Cord Injury (SCI) (Anderson, 2004; Ang et al., 2015) and the first rehabilitation target after stroke (Canning et al., 1999; Ada et al., 2006). Mainly, force generation is essential during our physical interactions (Westling and Johansson, 1984; Ostry and Feldman, 2003) and force control allows us to stabilise our manipulations (Franklin et al., 2003), making them safe and precise (Faisal et al., 2010; Branco et al., 2019b).

As we reviewed in Chapter 2, the opportunities that force generation offers in motor control restoration and rehabilitation have translated in remarkable progress in the invasive BCI direction, bringing as closer to understand single-neuron force encoding (Phillips, 1977; Cheney and Fetz, 1980; Porter and Lemon, 1993), force decoding from neuronal-ensembles (Carmena et al., 2003; Chhatbar and Francis, 2013) and their practical application in human force restoration (Bouton et al., 2016; Ganzer et al., 2020). While invasive approaches offer fundamental insights on force control mechanisms and their restoration, they fail to reach most of their potential users due to the cost and risks of the surgeries

involved (Wolpaw et al., 2006; McGie et al., 2013; Klein, 2016; Burwell et al., 2017).

In contrast, non-invasive BCIs are cheaper and safer to use and do not require surgeries (Wolpaw et al., 2006). As a consequence, these systems are more suitable for long-term use since they do not cause neuroglial scarring (McGie et al., 2013; Klein, 2016), and for short-term ambulatory use and rehabilitation since they can be easily worn and taken off (Silvoni et al., 2011; Kaiser et al., 2011; Bamdad et al., 2015).

Nonetheless, non-invasive systems have lower SNR than their invasive counterpart (Ball et al., 2009; Buzsáki et al., 2012). As we identified in our review, fNIRS and EEG, two popular non-invasive brain-sensing techniques, have been used together in an attempt to increase the SNR (e.g. Wolpaw and McFarland, 2004; Hansen et al., 2014; Yin et al., 2015; Khan and Hong, 2017) but two challenges stand on the way of reaching their full force decoding potential. First, the use of fNIRS by itself has been limited due to the lack of understanding between the Haemodynamic Responses (HRs) it measures and the generated force (Derosière et al., 2014; Shibuya et al., 2016; Wriessnegger et al., 2017) and a majority of methods being based on incomplete fixed models of that relationship (Abdelnour and Huppert, 2009; Ye et al., 2009; Jang et al., 2009) or requiring additional instrumentation (Zhang et al., 2009; Gagnon et al., 2011). Second, EEG and fNIRS measure very different physiological phenomena, mainly reflecting in very different signal time-scales (Hong et al., 2018), requiring an extensive exploration of the features that can offer the best combination of those modalities (Li et al., 2017; Khan and Hasan, 2020) and a lack of methods to combine them efficiently (Ahn and Jun, 2017; Hong et al., 2018).

Interestingly, these two aspects emerge in the context of recent invasive evidence suggesting that the cortex modulates the generated force in a non-linear fashion (Ethier et al., 2016; Naufel et al., 2019). In this situation, we identified Deep Learning (DL) —a non-linear automatic feature extractor technique (Bengio, 2012; LeCun et al., 2015)— as a suitable framework to process and interpret multimodal non-invasive brain-signals, which we develop and demonstrate in this thesis for the decoding of force.

Consequently, we generated a first multimodal (fNIRS, EEG, force, EMG, EOG and breathing) dataset reflecting on real-world force generation characteristics, HYGRiP (Chapter 3), that we release to the public to stimulate the development of fNIRS-EEG fusion algorithms (Ortega et al., 2020). In particular, previous research has mainly focused on slow paced finger pinches (Nambu et al., 2009) or force imageries (Coyle et al., 2004; Sitaram et al., 2007) to show lateralisation effects that disappear

when whole hand grasps are executed (Shibuya et al., 2008; Derosière et al., 2014; Shibuya et al., 2016; Wriessnegger et al., 2017). Since the velocity (Obrig et al., 1997; Kuboyama et al., 2005; Leff et al., 2011) and the intensity (Shibuya et al., 2008; Shibusawa et al., 2009) of hand contractions had already been linked to increased variability in the HRs detected by fNIRS, we designed a strenuous force task with high velocity and intensity contractions. We validated the consistency and synchronisation of our data and confirmed the high variability of the HRs, which we suggest is the result of the stimulation being faster than the auto-regulation time constant of the HR as in fMRI (Glover, 1999).

Following the inconclusive fNIRS hemispherical lateralisation results during force generation (Shibuya et al., 2008; Derosière et al., 2014; Shibuya et al., 2016; Wriessnegger et al., 2017), we benchmarked conventional methods and our own DL fNIRS processing method, hemCNN, using the fNIRS data in HYGRiP (Chapter 4). In particular, our hemCNN method outperformed conventional linear methods using fixed HR descriptors (Ye et al., 2009; Derosière et al., 2014) thanks to the CNN layers' ability to optimise their non-linear filters to HRs evoked during strenuous force generation.

Remarkably, hemCNN was better at generalising across subjects, which is typically avoided in BCI decoders (Nambu et al., 2009; Hiwa et al., 2016; Trakoolwilaiwan et al., 2017; Paek et al., 2019). Furthermore, by embedding the lateralisation hypothesis in our hemCNN, we could open the black-box model and interpret all the layers in our architecture, which other CNN methods do not allow (Hiwa et al., 2016; Trakoolwilaiwan et al., 2017). Interestingly, hemCNN learned to detect bifurcations in HbO and HbR as the signatures of cortical activation, allowing for accurate cross-subject detection. HbR to HbO subtraction (Yin et al., 2015) or separation (Khan et al., 2018) and early dips in the HbR concentration (Li et al., 2017; Zafar et al., 2019) had been previously suggested as features that improved fNIRS classification accuracy. However, these kind of engineered features have been found to be very sensitive to relative delays between HbO and HbR changes (Boden et al., 2007) influenced by systemic Hb variations (Kirilina et al., 2012). Indeed, this sensitivity to delays also led our standard LIs analysis to be inconclusive in our HYGRiP data (Chapter 4). In contrast, our hemCNN approach learned filters that were sensitive to the HbO-HbR delay and led to the highest activation in the range of functional neurocoupling delays ($1.6 \pm 0.2\,\mathrm{s}$; Boden et al., 2007).

Ultimately, the ability of hemCNN to find behaviourally driven and interpretable features allowed us to find early hand-specific bi-hemispheric activation patterns during force generation that were only observed on slower and longer stimulations ($\geqslant 30\,\mathrm{s}$; Shibuya et al., 2008, 2016) or were previously interpreted as global symmetrical activations (Derosière et al., 2014; Wriessnegger et al., 2017) in hand-grip

tasks. In contrast, hemCNN revealed that even though both hemispheres are soon (1.5 – 9 s) simultaneously active, each hand has a specific activation topography on each hemisphere, more evident around the 10-20 central (C3 and C4) positions above the hands' nodes. Previous evidence supports this kind of bihemispheric hand-specific spatial organisation on simpler finger-pinch tasks with slower time resolution (12 s; Nambu et al., 2009) or using invasive approaches (Serrien et al., 2006; Ames and Churchland, 2019). This thesis reveals this kind of bi-hemispheric engagement with faster time resolution and during more strenuous tasks. HemCNN's full interpretability and adaptive signal processing capabilities advance the field of machine learning in neurotechnology and neuroimaging. These capabilities can have the most significant impact in the analysis of altered haemodynamic function, e.g. after stroke or concussion, that render HR models (e.g. Buxton et al., 1998, 2004) inaccurate (Pineiro et al., 2002; Bonakdarpour et al., 2007; Stilling et al., 2019).

While the study of unimanual forces offers essential insights into the relationship of cortical activities with end-point forces, most skilled real-world manipulations require the simultaneous use of both hands (Faisal et al., 2010; Lima et al., 2020; Haar and Faisal, 2020). In this sense, from a BCI perspective, people requiring assistance, rehabilitation or restoration of one or both hands need hand-specific decoders that can discriminate the activities associated with each hand, avoiding undesired hand cross-talk effects that would hinder the coordination of the decoded outputs. This kind of control amounts to enable 2 (one per hand) simultaneous Degrees of Freedom (DoF) in BCI operation. However, as our systematic review identified (Chapter 2), most force decoding research is mostly focused in unimanual decoding (e.g. Nambu et al., 2009; Paek et al., 2019; Naufel et al., 2019; Ganzer et al., 2020). Moreover, the non-invasive BCI approaches that currently allow to operate different DoFs do so in sequence (Wolpaw and McFarland, 2004; LaFleur et al., 2013; Khan and Hong, 2017). However, since the user needs to wait to finish operating in one dimension before moving on to the next, this kind of sequential DoF control translates into longer operation periods.

In Chapter 5, we generated a second multimodal dataset with the particular goal of stimulating the study of non-invasive bimanual simultaneous control and the exploration of BCI algorithms with 2 simultaneous DoF decoding capabilities. Namely, HYGRiiP's task required both hands to produce different graded forces at the same time. This hand-specific engagement in the task resulted in interesting correlation structures between the hands and the fNIRS signals, being weaker for EEG (HYGRiiP's validation, Chapter 5). This kind of correlation builds on HYGRiP's observation that fast hand-grips ($\gg$ 1 Hz) do not allow cortical neurons enough time to resynchronise before a new desynchronisation

occurs, which might lead to weaker correlations between dynamic fast force onsets and EEG mu-ERDs (Chapter 3; Pfurtscheller et al., 2006). However, these correlations are just first-order approximations that might be unable to capture more complex dependencies between EEG and force generation, as previously suggested (Paek et al., 2019). In this context, HYGRiiP opens the exciting question of 2 simultaneous DoF control in non-invasive BCI and offers new opportunities in studying the modulation of brain signals with force building beyond first-order linear approaches, building on previous invasive research that suggests that 2 simultaneous DoF is possible (Ethier et al., 2016), and the modulation is non-linear (Ethier et al., 2016; Naufel et al., 2019).

In Chapter 6 we embarked on HYGRiiP's challenges in order to bring surgery-free non-invasive BCIs closer to the state-of-the-art invasive performance. To improve the lower non-invasive signals SNR (Ball et al., 2009; Buzsáki et al., 2012), we set out to find efficient ways of combining them. We found that the linear fusion of EEG and fNIRS results in better decoding performance than either modality alone, expanding previous research (Hansen et al., 2014; Yin et al., 2015; Chiarelli et al., 2018; Ghonchi et al., 2020) to the continuous bimanual force reconstruction, which was only addressed before with EEG in the unimanual case (Paek et al., 2019). Moreover, the continuous reconstruction of bimanual forces can be improved with our proposed *cnnatt* deep-fusion architecture. Hence, we expand recent evidence showing that DL improves the classification of multimodal brain signals (Chiarelli et al., 2018; Ghonchi et al., 2020) to continuous bimanual forces answering to the need for efficient EEG-fNIRS fusion methods (Ahn and Jun, 2017; Hong et al., 2018) that do not require extensive feature engineering. Besides, beyond mere reconstruction performance, we provide different performance measures and tests to understand multimodal brain signals' deep-fusion better. While most of the reconstruction originates in the better ability of DL to detect the application of force, we find that this detection is sensitive of the hand used, which previous studies could not verify due to the lack of simultaneous bimanual tasks (Nambu et al., 2009; Paek et al., 2019). Additionally, we observe that the main advantage in using DL consists in the better exploitation of the rich temporal scales within EEG and across EEG and fNIRS. By interpreting of *cnnatt* deep-fusion encoding, we offer interesting neurotechnological insights showing that DL is a useful behaviourally driven neuroimaging tool that allows for the concurrent use of different brain signals modalities. Indeed, we found hand specific cortical representations of force generation even during simultaneous hand use, and identifying activity patterns associated with unimanual error correction (Wang et al., 2005; Cavanagh et al., 2009; Womelsdorf et al., 2010), unimanual top-down control (hemCNN, Chapter 4; Sadaghiani and Kleinschmidt, 2016; Iturrate et al., 2018; Paek et al., 2019) and bimanual coordination (Serrien et al., 2006; Ames and Churchland, 2019). Remarkably,

HYGRiiP's task allowed *cnnatt* to reveal first-time evidence of hand-specific traces of non-linear non-invasive brain signals modulation with force, The existence of such modulation is a longstanding question for which Paek et al. (2019) could only discard a linear relationship. We offer new analytic opportunities to resolve this question and suggest that the new generation of non-invasive multimodal brain senors (Torricelli et al., 2014; Phan et al., 2016; Tang et al., 2016; De Roever et al., 2017) in combination with our deep-fusion techniques will be closer to achieve the bimanual continuous decoding performance that BCI users need, exempt from surgeries or invasive implants.

In summary, this thesis offers answers to five fundamental questions. Namely, we have shown: (1) how fNIRS can be used to physiologically decode motor tasks that challenged conventional approaches when it is processed with appropriately designed data-driven techniques; (2) that multimodal non-invasive brain signals (EEG and fNIRS) provide complementary information that improves the decoding of continuous bimanual force; (3) that DL improves the fusion of brain signals with different physical origins in a time-step basis; (4) that hand-specific brain encoding can be revealed during simultaneous hand control processing non-invasive signals with DL; and finally, to a more limited extent (5) that traces of hand-force modulated brain signals can be captured non-invasively using EEG and fNIRS and non-linear techniques.

In answering these questions, we opened new venues in which we discussed how the field of non-invasive BCI decoders for motor control could move forward. Specifically, BCI should be able to decode continuous variables with two or more DoFs, helping to restore or assist bimanual manipulations at an appropriate functional level, enabling users to carry out daily activities. As per our results, we identified that a promising venue is the combination of multimodal brain-sensing technologies with DL decoders.

7.2 Immediate applications

The contents of this thesis have a range of immediate applications. First, we digested neuroscientific data and processing pipelines to invite more machine learning experts into the field of neuroscience and BCI. The HYGRiP dataset can be used to develop decoding algorithms for unimanual strenuous force decoding. The HYGRiiP dataset can be used to develop similar algorithms for bimanual dynamic force decoding. Since the datasets share the same brain signals and recording setup, both can be used to study the transfer of brain function knowledge representations across tasks and subjects (Völker et al.,

2018; Wei et al., 2021). Namely, in both datasets, force is generated in an isometric way. In one dataset with each hand in sequence, and in the other with both hands simultaneously. We showed that the brain activity elicited during bimanual control is not only the sum of both hands unimanual control. Nonetheless, some of the brain function may still be dedicated to each hand's individual control even while they are used simultaneously (Ames and Churchland, 2019), which might be exploited to develop more sophisticated decoding techniques. In this sense, the datasets in conjunction directly offer an exciting test-bed to develop algorithms that find brain function signatures that can be transferred across force generation tasks, subjects and across hands.

HemCNN (Chapter 4) was developed with different fNIRS decoding and processing purposes in mind. In particular, the interpretable design makes it suited to study HRs in cases where the current white or grey-box models do not offer an appropriate explanation of the observed data. This was the case of our strenuous force generation task, for which conventional methods or studies using similar tasks were unable to find fNIRS hand-specific cortical signatures. Similarly, it is expected that the HR would not fit the available white and grey-box models in conditions such as internal bleeding, concussion or stroke (Pineiro et al., 2002; Bonakdarpour et al., 2007). We anticipate that hemCNN can be applied in these and similar cases to differentiate healthy from unhealthy HRs expanding on previous machine learning applications for the diagnosis and monitoring of brain diseases (Subhani et al., 2017; Gemein et al., 2020; Cicalese et al., 2020).

HemCNN can also be applied in hand force rehabilitation therapies (Dobkin, 2005; Xu et al., 2017; Baniqued et al., 2021). In this case, the hemCNN decoding signal could offer feedback to a user who is retraining to generate force with the left and the right hand expanding on current neurofeedback rehabilitation settings (Renton et al., 2017; Wang et al., 2018). The advantage of using hemCNN is that it might provide a rehabilitation signal that targets the recovery at the cortical level. In this case, changes in the cortex following rehabilitation may even appear before any behavioural recovery could be perceived. This capability would be of great value to learn what BCI interventions are better at providing sooner cortical function rehabilitation or detect changes at the cortical level that lead to motor function improvements. The early detection of cortical recovery could help to keep patients engaged and to develop user-adapted and evidence-based rehabilitation therapies (Winstein and Varghese, 2018).

Chapter 6 directly deals with the continuous control of force with 2-Degrees of Freedom (DoF) in the BCI context. Methods like *cnnatt*, or inspired in our work, will eventually unlock the continuous

control of the forces of two hands actuated simultaneously but independently of each other in BCIs. This simultaneous and independent hand activity decoding will immediately translate on a broader range of interactions and manipulations that BCI users will be able to perform. Similarly, these interactions will be more natural since they will be continuous.

As for hemCNN, these results can have implications not only in the recovery of force control but also in rehabilitation, which currently lacks these type of bimanual interventions (Kantak et al., 2017; Itkonen et al., 2019). If the cortical deterioration is not severe and allows patients to generate forces, they can use the methods we developed and demonstrated to control an assistive robotic device (Itkonen et al., 2019). Otherwise, subjects with a broken cortico-muscular neural link, would require a generalisable relationship between force and cortical signals since the force signal could not be acquired directly. In the most direct case, *cnnatt* or *cnnatt-lstm* can be used to build robotic hand-grip decoders (Petković et al., 2012) to assist and potentially rehabilitate hand-grips in subjects who can still perform force control but who do not perform well enough due to some neurological condition (Rayegani et al., 2014; Renton et al., 2017; Wang et al., 2018).

Finally, both hemCNN and *cnnatt*, can be directly applied to neuroimaging and be used as behavioural-data-driven imaging techniques following bottom-to-top approaches. These approaches could automatically find the cortical encoding of normal or altered (or normal vs altered) behaviours, currently limited to statistical differences in top-to-bottom approaches (Phillipou et al., 2018), and would benefit from the interpretability of our approaches (Davatzikos, 2019; Jollans et al., 2019; Kohoutová et al., 2020).

7.3 Future work

Due to our work's multidisciplinary nature, future directions can be better considered on three levels bearing in mind that these are very interrelated.

Future work in machine learning At the machine learning level, one of the most straightforward ways to move forward would be to collect a higher number of trials per subject. However, this will incur a time expense that not many researchers can afford. Instead, the existing data in HYGRiP and HYGRiiP can be used to investigate ways to transfer what the models learn in some subjects to the rest. This cross-data knowledge transfer can be achieved with transfer-learning and meta-learning techniques (Pan and Yang, 2009; Bengio, 2012; Völker et al., 2018; Wei et al., 2021). In this sense,

due to its interpretability and lightweight, we believe hemCNN and HYGRiP provide an excellent test-bed for developing and testing algorithms that represent brain function and analysing how knowledge representations transfer from one subject to the other. We already offer a baseline for cross-subject testing, which could be further improved with these transfer learning techniques.

Another attractive venue at the machine learning level is to enable methods to process inputs that present sparse eventful information in time . For example, signals that sometimes are very informative about an event but the rest of the time are noise. As we discussed in Chapter 6, brain signals (in particular EEG) carry most of their information in the form of events that happen sparsely in time (Vidal, 1973, 1977). The lack of alignment of such events can hinder their detection across channels and modalities. To some extent, we improved the detection of these sparse events using a self-attention module in our architecture that allows for their alignment (Bahdanau et al., 2014; Vaswani et al., 2017). While the previous approaches achieve outstanding performance in Natural Language Processing (NLP), most of the elements in natural language sequences (i.e. words) have a specific meaning or add nuances (Chowdhury, 2003). In contrast, as previously mentioned, this is not the case for non-invasive brain-signals for which informative events could be separated by even several seconds of non-informative sequences, particularly in EEG (Vidal, 1973, 1977). In this case, embedding attention into recurrences so that the model can retain informative event across noisy signal patches could offer new opportunities in aligning events across channels, modalities and also long time sequences.

One possible way to enable this would be to embed uncertainty in the decoded outputs, i.e. all non-informative time-steps in the input would output high uncertainty (Sensoy et al., 2018, 2020). We hypothesise that adding this extra uncertainty dimension would allow the attention mechanism to focus on aligning cross-modality events that are close in time, and the recurrence, to learn long-term dependencies on those locally cross-modality aligned events, further improving our results (*cnnatt-lstm*). Hence, the architecture would focus more on learning the most informative events, while the remaining are uncertain.

This mechanism could be enabled at learning time using semi-supervised (Wulsin et al., 2011; Jia et al., 2014) or meta-learning (Li et al., 2021; Duan et al., 2020) techniques that allow the architecture to decide what parts of a continuous output, aligned with noisy inputs, should not be taken into account. Here, gradients would need to propagate only through the time steps that are automatically detected to be informative, skipping the noisy ones. This skipping mechanism acquires particular importance

in the multimodal input case, as different inputs (e.g. fNIRS and EEG) can show different times in which they are informative or not, complementing each other. By enabling the uncertainty mechanism, a Backpropagation Through Time (BTT) training scheme could also learn flexible and independent memory depths for each modality without needing to fix them in advance.

Although not included here, we started exploring this venue building from work in synthetic handwriting (Graves, 2013; Chung et al., 2017) and uncertainty in neural networks (Sensoy et al., 2018, 2020). We started observing a positive effect of introducing edges of temporal informative events in the architecture. However, these recurrent architectures are custom made and do not include most of the optimisations that exist nowadays for off-the-shelf LSTM and similar recurrent architectures in packages like pytorch and tensorflow, resulting in longer training times.

Future work in neuroscience and BCIs At the neuroscientific level, one interesting question arising from our work, which could also be enabled by transfer-learning, is the possibility to transfer one hand control to the other, which has been explored before in fMRI (Haar et al., 2017). This cross-hand learning would enable people with only one healthy limb to regain control of a non-functional limb or prosthesis by training an algorithm on the healthy limb. However, to fully enable the missing limb decoding, the decoder should be made specific of the desired hand as we introduced in Chapters 5 and 6. The HYGRiP and HYGRiiP datasets can be used for this end. The idea is to find a functional relationship between the hand in a set of subjects and transfer it to an unseen subject for whom only one of the hands is used to adapt the decoder to the (virtual) unhealthy hand. The (virtual) unhealthy hand can be used to test the performance of the decoding. In this sense, using HYGRiiP in Chapter 6 we provide a benchmark on hand specific continuous decoding and ways to evaluate the decoding performance. Our results can be directly used as a starting point to develop this kind of hand-transfer decoding approaches for amputees, for whom plastic changes should be taken into consideration (Makin et al., 2015). A successful approach in this direction would open the door to start testing these algorithms in close-loop (only using the decoded signal to control a robotic arm or similar device to produce force) like recently demonstrated with invasive approaches (Ganzer et al., 2020).

Behaviourally driven neuroimaging Additionally, the techniques we proposed here could be used to further analyse neuroimaging data during unimanual and bimanual force production in fMRI studies

that are affected by similar sources of noise than fNIRS (Heinzel et al., 2013; Tong et al., 2019). Furthermore, the work presented in this thesis could help study brain function in unconstrained real-world tasks (Lima et al., 2020; Haar and Faisal, 2020) by allowing the recorded behaviour to drive the automatic learning of the most relevant cortical activation features.

Future work in rehabilitation with BCIs More directly, the models delivered here can be applied to the rehabilitation of bimanual control after stroke or brain injury (Dobkin, 2005; Xu et al., 2017; Baniqued et al., 2021). We provide trained deep models for fNIRS discrete decoding (hemCNN) and multimodal bimanual continuous force decoding (*cnnatt*). These models can be used to develop recovery measures during rehabilitation or create BCIs that assist in rehabilitation exercises using robots. The most immediate way to achieve this would be to use all the trained models provided as an ensemble (Dietterich, 2000; Zhou, 2012). In this way, the model ensemble would provide a sample of decoded outputs that can better characterise their uncertainty, obtaining a more robust way to create those rehabilitation indicators. Additionally, since hemCNN has straightforward interpretation in all layers, these new deep features can be explored as new neural markers of hemodynamic activity during rehabilitation expanding on previous hand-crafted features for this purpose (Venkatakrishnan et al., 2014) or use them to direct non-invasive brain stimulation during rehabilitation (Takeuchi and Izumi, 2013; Plow et al., 2013) expanding the use of our decoders to the closed-loop case.

Future work in BCI sensors Finally, we believe that new technological advances will soon bring new opportunities for the research presented here to reach its full potential.

Notably, in contrast to EEG, photonic systems, like fNIRS are yet to reach their physical limit. Indeed, technical advances keep pushing the limits of non-invasive brain-sensing forward. For example, fNIRS can obtain tissue depth-dependent optical properties and resolve oxygen saturation indices in the cortex (Torricelli et al., 2014; Kassab et al., 2018; Lacerenza et al., 2020) and a portable device will soon be commercialised (Kernel, 2020). Oxygen saturation is an absolute measure of the amount of oxygen in the blood and is used for diagnostic and critical care (Lange and Tachtsidis, 2019). The availability of this measure in the brain would allow us to study brain oxygenation differences across activities and diseases, defining ranges of normality and potentially validating fNIRS as a clinical tool.

Photonics research has also focused on detecting more specific and faster neuronal activities by targeting molecular activity in the respiratory chain (Phan et al., 2016; Nosrati et al., 2016; De Roever et al.,

2017; Lange et al., 2018). The respiratory chain takes places in the mitochondria and is in charge of producing the cells' metabolic energy. Targeting the respiratory chain activity offers a direct measure of cellular activity compared to the slow and indirect changes in blood oxygenation after neuronal metabolic demands (De Roever et al., 2017; Lange et al., 2018). Other approaches combine ultrasound and photonics (i.e. photoacoustics) and can provide anatomical information apart from haemoglobin concentration (Xia et al., 2014; Tang et al., 2016; Zhang et al., 2018). Photoacoustics increase the measure's spatio-temporal resolution and have been successfully applied in animal models (Tang et al., 2016; Zhang et al., 2018). Soon these technical advances might translate into more robust and accurate non-invasive BCI decoders.

Furthermore, more specific neural markers and physical multimodal interactions are being researched in the photonics field (Torricelli et al., 2014; Phan et al., 2016; De Roever et al., 2017). As it happened with fNIRS (Ferrari and Quaresima, 2012) and EEG (Casson, 2019), it might be just a matter of time that these new physical modalities are engineered to offer more stability and repeatability in smaller devices (Tang et al., 2016). The methods presented here could be easily applied to these new measures to exploit the faster time-scales and neural specificity that these new physical sensing modalities promise. Since the methods presented here are data-driven and agnostic of the inputs' physical origin, they could be applied to any physical modality without difficulty. Besides, we showed that DL is a valuable tool for the fusion of different measurement modalities for decoding continuous motor behaviour on a time-step basis. We suggest that the techniques we presented will be very valuable in exploiting novel non-invasive brain-sensing multimodal techniques accelerating the development of more robust and versatile BCIs.

These research directions, and others we cannot yet imagine, will soon be a reality for motivated researchers to embark on. We here provide valuable materials, tools and insights that advance the field of non-invasive and surgery-free BCI. With these, we contribute to a collective effort that ultimately aims to improve the autonomy and quality of life of people with disabilities and the development of accurate and cheaper tools for neurological diagnostics and monitoring.

Bibliography

Abdelnour, A. F. and Huppert, T. (2009). Real-time imaging of human brain function by near-infrared spectroscopy using an adaptive general linear model. *Neuroimage*, 46(1):133–143.

Ada, L., Dorsch, S., and Canning, C. G. (2006). Strengthening interventions increase strength and improve activity after stroke: a systematic review. *Australian Journal of Physiotherapy*, 52(4):241–248.

Aggarwal, V., Mollazadeh, M., Davidson, A. G., Schieber, M. H., and Thakor, N. V. (2013). State-based decoding of hand and finger kinematics using neuronal ensemble and LFP activity during dexterous reach-to-grasp movements. *Journal of Neurophysiology*, 109(12):3067–3081.

Ahn, S. and Jun, S. C. (2017). Multi-modal integration of EEG-fNIRS for brain-computer interfaces–current limitations and future directions. *Frontiers in Human Neuroscience*, 11:503.

Alexander, D. M., Jurica, P., Trengove, C., Nikolaev, A. R., Gepshtein, S., Zvyagintsev, M., Mathiak, K., Schulze-Bonhage, A., Ruescher, J., Ball, T., et al. (2013). Traveling waves and trial averaging: the nature of single-trial and averaged brain responses in large-scale cortical signals. *Neuroimage*, 73:95–112.

Ames, K. C. and Churchland, M. M. (2019). Motor cortex signals corresponding to the two arms are shared across hemispheres, mixed among neurons, yet partitioned within the population response. *bioRxiv*, pages 1–45.

Andersen, R. A. and Cui, H. (2009). Intention, action planning, and decision making in parietal-frontal circuits. *Neuron*, 63(5):568–583.

Anderson, K. D. (2004). Targeting recovery: priorities of the spinal cord-injured population. *Journal of Neurotrauma*, 21(10):1371–1383.

Ang, K. K., Chua, K. S. G., Phua, K. S., Wang, C., Chin, Z. Y., Kuah, C. W. K., Low, W., and Guan, C. (2015). A randomized controlled trial of EEG-based motor imagery brain-computer interface robotic rehabilitation for stroke. *Clinical EEG and Neuroscience*, 46(4):310–320.

Anumanchipalli, G. K., Chartier, J., and Chang, E. F. (2019). Speech synthesis from neural decoding of spoken sentences. *Nature*, 568(7753):493–498.

Arthur, K. C., Calvo, A., Price, T. R., Geiger, J. T., Chio, A., and Traynor, B. J. (2016). Projected increase in amyotrophic lateral sclerosis from 2015 to 2040. *Nature Communications*, 7(1):1–6.

Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K. R., and Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS One*, 10(7):1–46.

Bahdanau, D., Cho, K., and Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Ball, T., Demandt, E., Mutschler, I., Neitzel, E., Mehring, C., Vogt, K., Aertsen, A., and Schulze-Bonhage, A. (2008). Movement related activity in the high gamma range of the human EEG. *Neuroimage*, 41(2):302–310.

Ball, T., Kern, M., Mutschler, I., Aertsen, A., and Schulze-Bonhage, A. (2009). Signal quality of simultaneously recorded invasive and non-invasive EEG. *Neuroimage*, 46(3):708–716.

Ball, T., Schreiber, A., Feige, B., Wagner, M., Lücking, C. H., and Kristeva-Feige, R. (1999). The role of higher-order motor areas in voluntary movement as revealed by high-resolution EEG and fMRI. *Neuroimage*, 10(6):682–694.

Bamdad, M., Zarshenas, H., and Auais, M. A. (2015). Application of BCI systems in neurorehabilitation: a scoping review. *Disability and Rehabilitation: Assistive Technology*, 10(5):355–364.

Baniqued, P. D. E., Stanyer, E. C., Awais, M., Alazmani, A., Jackson, A. E., Mon-Williams, M. A., Mushtaq, F., and Holt, R. J. (2021). Brain–computer interface robotics for hand rehabilitation after stroke: a systematic review. *Journal of Neuroengineering and Rehabilitation*, 18(1):1–25.

Bansal, A. K., Truccolo, W., Vargas-Irwin, C. E., and Donoghue, J. P. (2012). Decoding 3D reach and grasp from hybrid signals in motor and premotor cortices: spikes, multiunit activity, and local field potentials. *Journal of Neurophysiology*, 107(5):1337–1355.

Belić, J. J. and Faisal, A. A. (2015). Decoding of human hand actions to handle missing limbs in neuroprosthetics. *Frontiers in Computational Neuroscience*, 9:27.

Bengio, Y. (2012). Deep learning of representations for unsupervised and transfer learning. In *Proceedings of ICML Workshop on Unsupervised and Transfer Learning*, pages 17–36.

Beni, N. H., Foodeh, R., Shalchyan, V., and Daliri, M. R. (2020). Force decoding using local field potentials in primary motor cortex: PLS or Kalman filter regression? *Physical and Engineering Sciences in Medicine*, 43(1):175–186.

Bocquelet, F., Hueber, T., Girin, L., Savariaux, C., and Yvert, B. (2016). Real-time control of an articulatory-based speech synthesizer for brain computer interfaces. *PLoS Computational Biology*, 12(11).

Boden, S., Obrig, H., Köhncke, C., Benav, H., Koch, S., and Steinbrink, J. (2007). The oxygenation response to functional stimulation: is there a physiological meaning to the lag between parameters? *Neuroimage*, 36(1):100–107.

Bonakdarpour, B., Parrish, T. B., and Thompson, C. K. (2007). Hemodynamic response function in patients with stroke-induced aphasia: implications for fMRI data analysis. *Neuroimage*, 36(2):322–331.

Bouton, C. E., Shaikhouni, A., Annetta, N. V., Bockbrader, M. A., Friedenberg, D. A., Nielson, D. M., Sharma, G., Sederberg, P. B., Glenn, B. C., Mysiw, W. J., Morgan, A. G., Deogaonkar, M., and Rezai, A. R. (2016). Restoring cortical control of functional movement in a human with quadriplegia. *Nature*, 533(7602):247–250.

Branco, M. P., de Boer, L. M., Ramsey, N. F., and Vansteensel, M. J. (2019a). Encoding of kinetic and kinematic movement parameters in the sensorimotor cortex: A Brain-Computer Interface perspective. *European Journal of Neuroscience*, 50(5):2755–2772.

Branco, M. P., Geukes, S. H., Aarnoutse, E. J., Vansteensel, M. J., Freudenburg, Z. V., and Ramsey, N. F. (2019b). High-frequency band temporal dynamics in response to a grasp force task. *Journal of Neural Engineering*, 16(5).

Burke, R., Levine, D., Tsairis, P., and Zajac Iii, F. (1973). Physiological types and histochemical profiles in motor units of the cat gastrocnemius. *The Journal of Physiology*, 234(3):723–748.

Burwell, S., Sample, M., and Racine, E. (2017). Ethical aspects of brain computer interfaces: a scoping review. *BMC Medical Ethics*, 18(1):1–11.

Buxton, R. B., Uludağ, K., Dubowitz, D. J., and Liu, T. T. (2004). Modeling the hemodynamic response to brain activation. *Neuroimage*, 23:S220–S233.

Buxton, R. B., Wong, E. C., and Frank, L. R. (1998). Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. *Magnetic Resonance in Medicine*, 39(6):855–864.

Buzsáki, G., Anastassiou, C. A., and Koch, C. (2012). The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6):407–420.

Buzsáki, G. and Wang, X.-J. (2012). Mechanisms of gamma oscillations. *Annual Review of Neuroscience*, 35:203–225.

Calautti, C. and Baron, J.-C. (2003). Functional neuroimaging studies of motor recovery after stroke in adults: a review. *Stroke*, 34(6):1553–1566.

Caldwell, M., Scholkmann, F., Wolf, U., Wolf, M., Elwell, C., and Tachtsidis, I. (2016). Modelling confounding effects from extracerebral contamination and systemic factors on functional near-infrared spectroscopy. *Neuroimage*, 143:91–105.

Canning, C. G., Ada, L., and O'Dwyer, N. (1999). Slowness to develop force contributes to weakness after stroke. *Archives of Physical Medicine and Rehabilitation*, 80(1):66–70.

Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., Patil, P. G., Henriquez, C. S., and Nicolelis, M. A. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biology*, 1(2):193–208.

Casson, A. J. (2019). Wearable EEG and beyond. *Biomedical Engineering Letters*, 9(1):53–71.

Casson, A. J., Yates, D. C., Smith, S. J., Duncan, J. S., and Rodriguez-Villegas, E. (2010). Wearable electroencephalography. *IEEE Engineering in Medicine and Biology Magazine*, 29(3):44–56.

Castiello, U. and Begliomini, C. (2008). The cortical control of visually guided grasping. *The Neuroscientist*, 14(2):157–170.

Cauraugh, J., Light, K., Kim, S., Thigpen, M., and Behrman, A. (2000). Chronic motor dysfunction after stroke: recovering wrist and finger extension by electromyography-triggered neuromuscular stimulation. *Stroke*, 31(6):1360–1364.

Cavanagh, J. F., Cohen, M. X., and Allen, J. J. (2009). Prelude to and resolution of an error: EEG phase synchrony reveals cognitive control dynamics during action monitoring. *Journal of Neuroscience*, 29(1):98–105.

Chen, C., Shin, D., Watanabe, H., Nakanishi, Y., Kambara, H., Yoshimura, N., Nambu, A., Isa, T., Nishimura, Y., and Koike, Y. (2014). Decoding grasp force profile from electrocorticography signals in non-human primate sensorimotor cortex. *Neuroscience Research*, 83:1–7.

Cheney, P. D. and Fetz, E. E. (1980). Functional classes of primate corticomotoneuronal cells and their relation to active force. *Journal of Neurophysiology*, 44(4):773–791.

Cheng, E. J., Brown, I. E., and Loeb, G. E. (2000). Virtual muscle: a computational approach to understanding the effects of muscle properties on motor control. *Journal of Neuroscience Methods*, 101(2):117–130.

Cherian, A., Krucoff, M. O., and Miller, L. E. (2011). Motor cortical prediction of EMG: Evidence that a kinetic brain-machine interface may be robust across altered movement dynamics. *Journal of Neurophysiology*, 106(2):564–575.

Chhatbar, P. Y. and Francis, J. T. (2013). Towards a naturalistic brain-machine interface: hybrid torque and position control allows generalization to novel dynamics. *PloS One*, 8(1):e52286.

Chiarelli, A. M., Croce, P., Merla, A., and Zappasodi, F. (2018). Deep learning for hybrid EEG-fNIRS brain–computer interface: application to motor imagery classification. *Journal of Neural Engineering*, 15(3):036028.

Chowdhury, G. G. (2003). Natural language processing. *Annual review of Information Science and Technology*, 37(1):51–89.

Chung, J., Ahn, S., and Bengio, Y. (2017). Hierarchical multiscale recurrent neural networks. In *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*.

Cicalese, P. A., Li, R., Ahmadi, M. B., Wang, C., Francis, J. T., Selvaraj, S., Schulz, P. E., and Zhang, Y. (2020). An EEG-fNIRS hybridization technique in the four-class classification of alzheimer's disease. *Journal of Neuroscience Methods*, 336:108618.

Cizek, V. (1970). Discrete Hilbert transform. *IEEE Transactions on Audio and Electroacoustics*, 18(4):340–343.

Cochin, S., Barthelemy, C., Roux, S., and Martineau, J. (1999). Observation and execution of movement: similarities demonstrated by quantified electroencephalography. *European Journal of Neuroscience*, 11(5):1839–1842.

Cope, M., Delpy, D., Reynolds, E., Wray, S., Wyatt, J., and Van der Zee, P. (1988). Methods of quantitating cerebral near infrared spectroscopy data. In *Oxygen Transport to Tissue X*, pages 183–189. Springer, New York.

Coyle, S., Ward, T., Markham, C., and McDarby, G. (2004). On the suitability of near-infrared (NIR) systems for next-generation brain–computer interfaces. *Physiological Measurement*, 25(4):815.

Croce, P., Zappasodi, F., Merla, A., and Chiarelli, A. M. (2017). Exploiting neurovascular coupling: a Bayesian sequential Monte Carlo approach applied to simulated EEG fNIRS data. *Journal of Neural Engineering*, 14(4):046029.

Cui, X., Bray, S., and Reiss, A. L. (2010). Functional near infrared spectroscopy (NIRS) signal improvement based on negative correlation between oxygenated and deoxygenated hemoglobin dynamics. *Neuroimage*, 49(4):3039–3046.

Custo, A., Wells Iii, W. M., Barnett, A. H., Hillman, E. M., and Boas, D. A. (2006). Effective scattering coefficient of the cerebral spinal fluid in adult head models for diffuse optical imaging. *Applied Optics*, 45(19):4747–4755.

Daly, J. J. and Wolpaw, J. R. (2008). Brain–computer interfaces in neurological rehabilitation. *The Lancet Neurology*, 7(11):1032–1043.

Daunizeau, J., Grova, C., Marrelec, G., Mattout, J., Jbabdi, S., Pélégrini-Issac, M., Lina, J.-M., and Benali, H. (2007). Symmetrical event-related EEG/fMRI information fusion in a variational Bayesian framework. *Neuroimage*, 36(1):69–87.

Davare, M., Duque, J., Vandermeeren, Y., Thonnard, J.-L., and Olivier, E. (2006). Role of the ipsilateral primary motor cortex in controlling the timing of hand muscle recruitment. *Cerebral Cortex*, 17(2):353–362.

Davare, M., Kraskov, A., Rothwell, J. C., and Lemon, R. N. (2011). Interactions between areas of the cortical grasping network. *Current Opinion in Neurobiology*, 21(4):565–570.

Davatzikos, C. (2019). Machine learning in neuroimaging: Progress and challenges. *Neuroimage*, 197:652.

De Roever, I., Bale, G., Cooper, R. J., and Tachtsidis, I. (2017). Functional NIRS measurement of Cytochrome-C-Oxidase demonstrates a more brain-specific marker of frontal lobe activation compared to the haemoglobins. In *Oxygen Transport to Tissue XXXIX*, pages 141–147. Springer.

Delpy, D. T., Cope, M., van der Zee, P., Arridge, S., Wray, S., and Wyatt, J. (1988). Estimation of optical pathlength through tissue from direct time of flight measurement. *Physics in Medicine & Biology*, 33(12):1433.

Derosière, G., Alexandre, F., Bourdillon, N., Mandrick, K., Ward, T. E., and Perrey, S. (2014). Similar scaling of contralateral and ipsilateral cortical responses during graded unimanual force generation. *Neuroimage*, 85:471–477.

Dettmers, C., Fink, G. R., Lemon, R. N., Stephan, K. M., Passingham, R. E., Silbersweig, D., Holmes, A., Ridding, M. C., Brooks, D. J., and Frackowiak, R. (1995). Relation between cerebral activity and force in the motor areas of the human brain. *Journal of Neurophysiology*, 74(2):802–815.

Di Bono, M. G., Begliomini, C., Castiello, U., and Zorzi, M. (2015). Probing the reaching–grasping network in humans through multivoxel pattern decoding. *Brain and Behavior*, 5(11):e00412.

Dietterich, T. G. (2000). Ensemble methods in machine learning. In *International Workshop on Multiple Classifier Systems*, pages 1–15. Springer.

Dimyan, M. A. and Cohen, L. G. (2011). Neuroplasticity in the context of motor rehabilitation after stroke. *Nature Reviews Neurology*, 7(2):76.

Dobkin, B. H. (2005). Rehabilitation after stroke. *New England Journal of Medicine*, 352(16):1677–1684.

Duan, T., Shaikh, M. A., Chauhan, M., Chu, J., Srihari, R. K., Pathak, A., and Srihari, S. N. (2020). Meta Learn on Constrained Transfer Learning for Low Resource Cross Subject EEG Classification. *IEEE Access*, 8:224791–224802.

Duncan, A., Meek, J. H., Clemence, M., Elwell, C. E., Tyszczuk, L., Cope, M., and Delpy, D. (1995). Optical pathlength measurements on adult head, calf and forearm and the head of the newborn infant using phase resolved optical spectroscopy. *Physics in Medicine & Biology*, 40(2):295.

Earl, R. and Hartley, N. (2019). A time for change. SIA Annual Report 2018/2019. spinal.co.uk. Accessed: Dec 2020.

Ehrsson, H. H., Fagergren, A., Jonsson, T., Westling, G., Johansson, R. S., and Forssberg, H. (2000). Cortical activity in precision-versus power-grip tasks: an fMRI study. *Journal of Neurophysiology*, 83(1):528–536.

Enriquez-Geppert, S., Huster, R. J., and Herrmann, C. S. (2013). Boosting brain functions: Improving executive functions with behavioral training, neurostimulation, and neurofeedback. *International Journal of Psychophysiology*, 88(1):1–16.

Ethier, C., Acuna, D., Solla, S. A., and Miller, L. E. (2016). Adaptive neuron-to-EMG decoder training for FES neuroprostheses. *Journal of Neural Engineering*, 13(4):046009.

Evarts, E. V. (2011). *Role of Motor Cortex in Voluntary Movements in Primates*, pages 1083–1120. American Cancer Society.

Fagg, A. H., Ojakangas, G. W., Miller, L. E., and Hatsopoulos, N. G. (2009). Kinetic trajectory decoding using motor cortical ensembles. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 17(5):487–496.

Faisal, A., Stout, D., Apel, J., and Bradley, B. (2010). The manipulative complexity of Lower Paleolithic stone toolmaking. *PloS One*, 5(11):e13718.

Faisal, A. A., Selen, L. P., and Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4):292.

Fatemi, M. and Daliri, M. R. (2020). Nonlinear sparse partial least squares: an investigation of the effect of nonlinearity and sparsity on the decoding of intracranial data. *Journal of Neural Engineering*, 17(1):016055.

Feigin, V. L., Abajobir, A. A., Abate, K. H., Abd-Allah, F., Abdulle, A. M., Abera, S. F., Abyu, G. Y., Ahmed, M. B., Aichour, A. N., Aichour, I., et al. (2017). Global, regional, and national burden of neurological disorders during 1990–2015: a systematic analysis for the Global Burden of Disease Study 2015. *The Lancet Neurology*, 16(11):877–897.

Ferrari, M. and Quaresima, V. (2012). A brief review on the history of human functional near-infrared spectroscopy (fNIRS) development and fields of application. *Neuroimage*, 63(2):921–935.

Figley, C. R. and Stroman, P. W. (2011). The role(s) of astrocytes and astrocyte activity in neurometabolism, neurovascular coupling, and the production of functional neuroimaging signals. *European Journal of Neuroscience*, 33(4):577–588.

Fischer, P., Pogosyan, A., Cheeran, B., Green, A. L., Aziz, T. Z., Hyam, J., Little, S., Foltynie, T., Limousin, P., Zrinzo, L., et al. (2017). Subthalamic nucleus beta and gamma activity is modulated depending on the level of imagined grip force. *Experimental Neurology*, 293:53–61.

Flint, R. D., Rosenow, J. M., Tate, M. C., and Slutzky, M. W. (2016). Continuous decoding of human grasp kinematics using epidural and subdural signals. *Journal of Neural Engineering*, 14(1):016005.

Flint, R. D., Wang, P. T., Wright, Z. A., King, C. E., Krucoff, M. O., Schuele, S. U., Rosenow, J. M., Hsu, F. P., Liu, C. Y., Lin, J. J., et al. (2014). Extracting kinetic information from human motor cortical signals. *Neuroimage*, 101:695–703.

Foodeh, R., Ebadollahi, S., and Daliri, M. R. (2020). Regularized Partial Least Square Regression for Continuous Decoding in Brain-Computer Interfaces. *Neuroinformatics*, pages 1–13.

Franklin, D. W., Osu, R., Burdet, E., Kawato, M., and Milner, T. E. (2003). Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. *Journal of Neurophysiology*, 90(5):3270–3282.

French, J. S., Anderson-Erisman, K. D., and Sutter, M. (2010). What do spinal cord injury consumers want? A review of spinal cord injury consumer priorities and neuroprosthesis from the 2008 neural interfaces conference. *Neuromodulation: Technology at the Neural Interface*, 13(3):229–231.

Fries, P. (2015). Rhythms for cognition: communication through coherence. *Neuron*, 88(1):220–235.

Friston, K. J., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. *Neuroimage*, 19(4):1273–1302.

Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., and Frackowiak, R. S. (1994). Statistical parametric maps in functional imaging: a general linear approach. *Human Brain Mapping*, 2(4):189–210.

Friston, K. J., Mechelli, A., Turner, R., and Price, C. J. (2000). Nonlinear responses in fMRI: the Balloon model, Volterra kernels, and other hemodynamics. *Neuroimage*, 12(4):466–477.

Fromm, C. (1983). Changes of steady state activity in motor cortex consistent with the length-tension relation of muscle. *Pflügers Archiv*, 398(4):318–323.

Fu, Y., Xiong, X., Jiang, C., Xu, B., Li, Y., and Li, H. (2016). Imagined hand clenching force and speed modulate brain activity and are classified by NIRS combined with EEG. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(9):1641–1652.

Gagnon, L., Cooper, R. J., Yücel, M. A., Perdue, K. L., Greve, D. N., and Boas, D. A. (2012). Short separation channel location impacts the performance of short channel regression in NIRS. *Neuroimage*, 59(3):2518–2528.

Gagnon, L., Perdue, K., Greve, D. N., Goldenholz, D., Kaskhedikar, G., and Boas, D. A. (2011). Improved recovery of the hemodynamic response in diffuse optical imaging using short optode separations and state-space modeling. *Neuroimage*, 56(3):1362–1371.

Ganesh, G., Burdet, E., Haruno, M., and Kawato, M. (2008). Sparse linear regression for reconstructing muscle activity from human cortical fMRI. *Neuroimage*, 42(4):1463–1472.

Ganzer, P. D., Colachis 4th, S. C., Schwemmer, M. A., Friedenberg, D. A., Dunlap, C. F., Swiftney, C. E., Jacobowitz, A. F., Weber, D. J., Bockbrader, M. A., and Sharma, G. (2020). Restoring the sense of touch using a sensorimotor demultiplexing neural interface. *Cell*, 181(4):763–773.

Gemein, L. A., Schirrmeister, R. T., Chrabąszcz, P., Wilson, D., Boedecker, J., Schulze-Bonhage, A., Hutter, F., and Ball, T. (2020). Machine-learning-based diagnostics of EEG pathology. *Neuroimage*, 220:117021.

Georgopoulos, A. P. and Ashe, J. (2000). One motor cortex, two different views. *Nature Neuroscience*, 3(10):963–963.

Ghonchi, H., Fateh, M., Abolghasemi, V., Ferdowsi, S., and Rezvani, M. (2020). Deep recurrent–convolutional neural network for classification of simultaneous EEG–fNIRS signals. *IET Signal Processing*, 14(3):142–153.

Glover, G. H. (1999). Deconvolution of impulse response in event-related BOLD fMRI1. *Neuroimage*, 9(4):416–429.

Gordon, A., Huxley, A. F., and Julian, F. (1966). The variation in isometric tension with sarcomere length in vertebrate muscle fibres. *The Journal of Physiology*, 184(1):170–192.

Graves, A. (2013). Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*.

Guenther, F. H., Brumberg, J. S., Wright, E. J., Nieto-Castanon, A., Tourville, J. A., Panko, M., Law, R., Siebert, S. A., Bartels, J. L., Andreasen, D. S., et al. (2009). A wireless brain-machine interface for real-time speech synthesis. *PloS One*, 4(12).

Gupta, R. and Ashe, J. (2009). Offline decoding of end-point forces using neural ensembles: application to a brain–machine interface. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 17(3):254–262.

Haar, S., Dinstein, I., Shelef, I., and Donchin, O. (2017). Effector-invariant movement encoding in the human motor system. *Journal of Neuroscience*, 37(37):9054–9063.

Haar, S. and Faisal, A. A. (2020). Brain activity reveals multiple motor-learning mechanisms in a real-world task. *Frontiers in Human Neuroscience*, 14:354.

Haas, L. F. (2003). Hans Berger (1873–1941), Richard Caton (1842–1926), and electroencephalography. *Journal of Neurology, Neurosurgery and Psychiatry*, 74(1):9–9.

Hansen, M. H., Kassebaum, E., Plocharska, M. A., Jochumsen, M., and Kamavuako, E. N. (2014). Improved detection and force decoding through combined near-infrared spectroscopy and electroencephalographic measurements. In *Replace, Repair, Restore, Relieve–Bridging Clinical and Engineering Solutions in Neurorehabilitation*, pages 411–419. Springer.

Hao, Y., Zhang, Q., Controzzi, M., Cipriani, C., Li, Y., Li, J., Zhang, S., Wang, Y., Chen, W., Carrozza, M. C., et al. (2014). Distinct neural patterns enable grasp types decoding in monkey dorsal premotor cortex. *Journal of Neural Engineering*, 11(6):066011.

Hayashi, M., Tsuchimoto, S., Mizuguchi, N., Miyatake, M., Kasuga, S., and Ushiba, J. (2019). Two-stage regression of high-density scalp electroencephalograms visualizes force regulation signaling during muscle contraction. *Journal of Neural Engineering*, 16(5):056020.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778.

Heffner, R. and Masterton, B. (1975). Variation in form of the pyramidal tract and its relationship to digital dexterity. *Brain, Behavior and Evolution*, 12(3):188–200.

Heinzel, S., Haeussinger, F. B., Hahn, T., Ehlis, A.-C., Plichta, M. M., and Fallgatter, A. J. (2013). Variability of (functional) hemodynamics as measured with simultaneous fNIRS and fMRI during intertemporal choice. *Neuroimage*, 71:125–134.

Hill, A. V. (1938). The heat of shortening and the dynamic constants of muscle. *Proceedings of the Royal Society of London. Series B - Biological Sciences*, 126(843):136–195.

Hirth, C., Obrig, H., Valdueza, J., Dirnagl, U., and Villringer, A. (1997). Simultaneous assessment of cerebral oxygenation and hemodynamics during a motor task. In Nemoto and LaManna, editors, *Oxygen Transport to Tissue XVIII*, pages 461–469. Springer, New York.

Hiwa, S., Hanawa, K., Tamura, R., Hachisuka, K., and Hiroyasu, T. (2016). Analyzing brain functions by subject classification of functional near-infrared spectroscopy data using convolutional neural networks analysis. *Computational Intelligence and Neuroscience*, 2016.

Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., Haddadin, S., Liu, J., Cash, S. S., Van Der Smagt, P., et al. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375.

Holper, L., Biallas, M., and Wolf, M. (2009). Task complexity relates to activation of cortical motor areas during uni- and bimanual performance: A functional NIRS study. *Neuroimage*, 46(4):1105–1113.

Hong, K.-S. and Khan, M. J. (2017). Hybrid brain–computer interface techniques for improved classification accuracy and increased number of commands: a review. *Frontiers in Neurorobotics*, 11:35.

Hong, K.-S., Khan, M. J., and Hong, M. J. (2018). Feature extraction and classification methods for hybrid fNIRS-EEG brain-computer interfaces. *Frontiers in Human Neuroscience*, 12:246.

Hoshi, Y. and Tamura, M. (1993). Dynamic multichannel near-infrared optical imaging of human brain activity. *Journal of Applied Physiology*, 75(4):1842–1846.

Huppert, T. J. (2016). Commentary on the statistical properties of noise and its implication on general linear models in functional near-infrared spectroscopy. *Neurophotonics*, 3(1):010401–010401.

Huppert, T. J., Hoge, R. D., Diamond, S. G., Franceschini, M. A., and Boas, D. A. (2006). A temporal comparison of BOLD, ASL, and NIRS hemodynamic responses to motor stimuli in adult humans. *Neuroimage*, 29(2):368–382.

Iso, N., Moriuchi, T., Sagari, A., Kitajima, E., Iso, F., Tanaka, K., Kikuchi, Y., Tabira, T., and Higashi, T. (2016). Monitoring local regional hemodynamic signal changes during motor execution and motor imagery using near-infrared spectroscopy. *Frontiers in Physiology*, 6:416.

Itkonen, M., Costa, Á., Yamasaki, H., Okajima, S., Alnajjar, F., Kumada, T., and Shimoda, S. (2019). Influence of bimanual exercise on muscle activation in post-stroke patients. *Robomech Journal*, 6(1):1–11.

Iturrate, I., Chavarriaga, R., Pereira, M., Zhang, H., Corbet, T., Leeb, R., and Millán, J. d. R. (2018). Human EEG reveals distinct neural correlates of power and precision grasping types. *Neuroimage*, 181(June):635–644.

Jang, K. E., Tak, S., Jung, J., Jang, J., Jeong, Y., and Ye, J. C. (2009). Wavelet minimum description length detrending for near-infrared spectroscopy. *Journal of Biomedical Optics*, 14(3):34004–34004.

Jasdzewski, G., Strangman, G., Wagner, J., Kwong, K., Poldrack, R., and Boas, D. (2003). Differences in the hemodynamic response to event-related motor and visual paradigms as measured by near-infrared spectroscopy. *Neuroimage*, 20(1):479–488.

Jia, X., Li, K., Li, X., and Zhang, A. (2014). A novel semi-supervised deep learning framework for affective state recognition on eeg signals. In *2014 IEEE International Conference on Bioinformatics and Bioengineering*, pages 30–37. IEEE.

Jochumsen, M., Niazi, I. K., Mrachacz-Kersting, N., Farina, D., and Dremstrup, K. (2013). Detection and classification of movement-related cortical potentials associated with task force and speed. *Journal of Neural Engineering*, 10(5):56015.

Jollans, L., Boyle, R., Artiges, E., Banaschewski, T., Desrivières, S., Grigis, A., Martinot, J.-L., Paus, T., Smolka, M. N., Walter, H., et al. (2019). Quantifying performance of machine learning methods for neuroimaging data. *Neuroimage*, 199:351–365.

Kaiser, V., Kreilinger, A., Müller-Putz, G. R., and Neuper, C. (2011). First steps toward a motor imagery based stroke BCI: new strategy to set up a classifier. *Frontiers in Neuroscience*, 5:86.

Kantak, S., Jax, S., and Wittenberg, G. (2017). Bimanual coordination: A missing piece of arm rehabilitation after stroke. *Restorative neurology and neuroscience*, 35(4):347–364.

Kassab, A., Le Lan, J., Tremblay, J., Vannasing, P., Dehbozorgi, M., Pouliot, P., Gallagher, A., Lesage, F., Sawan, M., and Nguyen, D. K. (2018). Multichannel wearable fNIRS-EEG system for long-term clinical monitoring. Technical report, Wiley Online Library.

Kato, T., Kamei, A., Takashima, S., and Ozaki, T. (1993). Human visual cortical function during photic stimulation monitoring by means of near-infrared spectroscopy. *Journal of Cerebral Blood Flow and Metabolism*, 13(3):516–520.

Keisker, B., Hepp-Reymond, M. C., Blickenstorfer, A., and Kollias, S. S. (2010). Differential representation of dynamic and static power grip force in the sensorimotor network. *European Journal of Neuroscience*, 31(8):1483–1491.

Kernel (2020). Flow50. www.kernel.com/flow. Culver city, Los Angeles (USA). Accessed: Dec 2020.

Khan, M. J., Ghafoor, U., and Hong, K.-S. (2018). Early detection of hemodynamic responses using EEG: a hybrid EEG-fNIRS study. *Frontiers in Human Neuroscience*, 12:479.

Khan, M. J. and Hong, K.-S. (2017). Hybrid EEG-fNIRS based eight-command decoding for BCI: application to quadcopter control. *Frontiers in Neurorobotics*, 11:6.

Khan, M. U. and Hasan, M. A. (2020). Hybrid EEG-fNIRS BCI fusion using Multi-Resolution Singular Value Decomposition (MSVD). *Frontiers in Human Neuroscience*, 14.

Khawaldeh, S., Tinkhauser, G., Shah, S. A., Peterman, K., Debove, I., Nguyen, T. K., Nowacki, A., Lachenmayer, M. L., Schuepbach, M., Pollo, C., et al. (2020). Subthalamic nucleus activity dynamics and limb movement prediction in Parkinson's disease. *Brain*, 143(2):582–596.

Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In *Proceedings of the 2015 3rd International Conference on Learning Representations*.

Kirilina, E., Jelzow, A., Heine, A., Niessing, M., Wabnitz, H., Brühl, R., Ittermann, B., Jacobs, A. M., and Tachtsidis, I. (2012). The physiological origin of task-evoked systemic artefacts in functional near infrared spectroscopy. *Neuroimage*, 61(1):70–81.

Klein, E. (2016). Informed consent in implantable BCI research: Identifying risks and exploring meaning. *Science and Engineering Ethics*, 22(5):1299–1317.

Klem, G. H., Lüders, H. O., Jasper, H., Elger, C., et al. (1999). The ten-twenty electrode system of the International Federation. *Electroencephalography and Clinical Neurophysiology*, 52(3):3–6.

Klimesch, W., Sauseng, P., and Hanslmayr, S. (2007). EEG alpha oscillations: the inhibition–timing hypothesis. *Brain Research Reviews*, 53(1):63–88.

Koh, P. H., Glaser, D., Flandin, G., Kiebel, S., Butterworth, B., Maki, A., Delpy, D. T., and Elwell, C. E. (2007). Functional optical signal analysis: a software tool for near-infrared spectroscopy data processing incorporating statistical parametric mapping. *Journal of Biomedical Optics*, 12(6):064010.

Kohoutová, L., Heo, J., Cha, S., Lee, S., Moon, T., Wager, T. D., and Woo, C.-W. (2020). Toward a unified framework for interpreting machine-learning models in neuroimaging. *Nature Protocols*, 15(4):1399–1435.

Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C., and Faisal, A. A. (2018). The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, 24(11):1716–1720.

Kristeva-Feige, R., Fritsch, C., Timmer, J., and Lücking, C.-H. (2002). Effects of attention and precision of exerted force on beta range EEG-EMG synchronization during a maintained motor contraction task. *Clinical Neurophysiology*, 113(1):124–131.

Kuboyama, N., Nabetani, T., Shibuya, K., Machida, K., and Ogaki, T. (2005). Relationship between cerebral activity and movement frequency of maximal finger tapping. *Journal of Physiological Anthropology and Applied Human Science*, 24(3):201–208.

Kwong, K. K., Belliveau, J. W., Chesler, D. A., Goldberg, I. E., Weisskoff, R. M., Poncelet, B. P., Kennedy, D. N., Hoppel, B. E., Cohen, M. S., and Turner, R. (1992). Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences*, 89(12):5675–5679.

Lacerenza, M., Buttafava, M., Renna, M., Dalla Mora, A., Spinelli, L., Zappa, F., Pifferi, A., Torricelli, A., Tosi, A., and Contini, D. (2020). Wearable and wireless time-domain near-infrared spectroscopy system for brain and muscle hemodynamic monitoring. *Biomedical Optics Express*, 11(10):5934–5949.

LaFleur, K., Cassady, K., Doud, A., Shades, K., Rogin, E., and He, B. (2013). Quadcopter control in three-dimensional space using a noninvasive motor imagery-based brain–computer interface. *Journal of Neural Engineering*, 10(4):046003.

Lange, F., Dunne, L., Hale, L., and Tachtsidis, I. (2018). MAESTROS: a multiwavelength time-domain NIRS system to monitor changes in oxygenation and oxidation state of Cytochrome-C-Oxidase. *IEEE Journal of Selected Topics in Quantum Electronics*, 25(1):1–12.

Lange, F. and Tachtsidis, I. (2019). Clinical brain monitoring with time domain NIRS: A review and future perspectives. *Applied Sciences*, 9(8):1612.

Lawrence, N. D. (2017). Data readiness levels. *arXiv preprint arXiv:1705.02245*.

Leahy, R., Mosher, J., Spencer, M., Huang, M., and Lewine, J. (1998). A study of dipole localization accuracy for MEG and EEG using a human skull phantom. *Electroencephalography and clinical neurophysiology*, 107(2):159–173.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.

Leff, D. R., Orihuela-Espina, F., Elwell, C. E., Athanasiou, T., Delpy, D. T., Darzi, A. W., and Yang, G.-Z. (2011). Assessment of the cerebral cortex during motor task behaviours in adults: a systematic review of functional near infrared spectroscopy (fNIRS) studies. *Neuroimage*, 54(4):2922–2936.

Li, D., Ortega, P., Wei, X., and Faisal, A. (2021). Meta-Learning EEG Motor Imagery Decoder for Brain Computer Interface. In *10th Interational IEEE EMBS Conference on Neural Engineering*. IEEE.

Li, R., Potter, T., Huang, W., and Zhang, Y. (2017). Enhancing performance of a hybrid EEG-fNIRS system using channel selection and early temporal features. *Frontiers in Human Neuroscience*, 11:462.

Lima, I. R., Haar, S., Di Grassi, L., and Faisal, A. A. (2020). Neurobehavioural signatures in race car driving: a case study. *Scientific Reports*, 10(1):1–9.

Maki, A., Yamashita, Y., Watanabe, E., and Koizumi, H. (1996). Visualizing human motor activity by using non-invasive optical topography. *Frontiers of Medical and Biological Engineering*, 7(4):285–297.

Makin, T. R., Scholz, J., Henderson Slater, D., Johansen-Berg, H., and Tracey, I. (2015). Reassessing cortical reorganization in the primary sensorimotor cortex following arm amputation. *Brain*, 138(8):2140–2146.

Mansouri, C., L'Huillier, J.-P., Kashou, N. H., and Humeau, A. (2010). Depth sensitivity analysis of functional near-infrared spectroscopy measurement using three-dimensional Monte Carlo modelling-based magnetic resonance imaging. *Lasers in Medical Science*, 25(3):431–438.

Matcher, S., Elwell, C., Cooper, C., Cope, M., and Delpy, D. (1995). Performance comparison of several published tissue near-infrared spectroscopy algorithms. *Analytical Biochemistry*, 227(1):54–68.

McFarland, D. J., Miner, L. A., Vaughan, T. M., and Wolpaw, J. R. (2000). Mu and beta rhythm topographies during motor imagery and actual movements. *Brain Topography*, 12(3):177–186.

McGie, S. C., Nagai, M. K., and Artinian-Shaheen, T. (2013). Clinical ethical concerns in the implantation of brain-machine interfaces: part I: overview, target populations, and alternatives. *IEEE Pulse*, 4(1):28–32.

McKiernan, B. J., Marcario, J. K., Karrer, J. H., and Cheney, P. D. (2000). Correlations between corticomotoneuronal (CM) cell postspike effects and cell-target muscle covariation. *Journal of Neurophysiology*, 83(1):99–115.

Meng, J., Zhang, S., Bekyo, A., Olsoe, J., Baxter, B., and He, B. (2016). Noninvasive electroencephalogram based control of a robotic arm for reach and grasp tasks. *Scientific Reports*, 6:38565.

Menon, R. S., Ogawa, S., Hu, X., Strupp, J. P., Anderson, P., and Uĝurbil, K. (1995). BOLD based functional MRI at 4 Tesla includes a capillary bed contribution: Echo-planar imaging correlates with previous optical imaging using intrinsic signals. *Magnetic Resonance in Medicine*, 33(3):453–459.

Miyai, I., Tanabe, H. C., Sase, I., Eda, H., Oda, I., Konishi, I., Tsunazawa, Y., Suzuki, T., Yanagida, T., and Kubota, K. (2001). Cortical mapping of gait in humans: a near-infrared spectroscopic topography study. *Neuroimage*, 14(5):1186–1192.

Mulliken, G. H., Musallam, S., and Andersen, R. A. (2008). Forward estimation of movement state in posterior parietal cortex. *Proceedings of the National Academy of Sciences*, 105(24):8170–8177.

Mussa-Ivaldi, F. (1988). Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neuroscience Letters*, 91(1):106–111.

Nakajima, K., Maier, M., Kirkwood, P., and Lemon, R. (2000). Striking differences in transmission of corticospinal excitation to upper limb motoneurons in two primate species. *Journal of Neurophysiology*, 84(2):698–709.

Nambu, I., Osu, R., Aki Sato, M., Ando, S., Kawato, M., and Naito, E. (2009). Single-trial reconstruction of finger-pinch forces from human motor-cortical activation measured by near-infrared spectroscopy (NIRS). *Neuroimage*, 47(2):628–637.

Naufel, S., Glaser, J. I., Kording, K. P., Perreault, E. J., and Miller, L. E. (2019). A muscle-activity-dependent gain between motor cortex and EMG. *Journal of Neurophysiology*, 121(1):61–73.

Nicolas-Alonso, L. F. and Gomez-Gil, J. (2012). Brain computer interfaces, a review. *Sensors*, 12(2):1211–1279.

Nomenclature, S. E. P. (1991). American Electroencephalographic Society guidelines for standard electrode position nomenclature. *Journal of Clinical Neurophysiology*, 8(2):200–202.

Nosrati, R., Vesely, K., Schweizer, T. A., and Toronov, V. (2016). Event-related changes of the prefrontal cortex oxygen delivery and metabolism during driving measured by hyperspectral fNIRS. *Biomedical Optics Express*, 7(4):1323–1335.

Obrig, H., Hirth, C., Junge-Hülsing, J. G., Döge, C., Wenzel, R., Wolf, T., Dirnagl, U., and Villringer, A. (1997). Length of resting period between stimulation cycles modulates hemodynamic response to a motor stimulus. In *Oxygen Transport to Tissue XVIII*, pages 471–480. Springer.

Okada, E. and Delpy, D. T. (2003). Near-infrared light propagation in an adult head model. II. Effect of superficial tissue thickness on the sensitivity of the near-infrared spectroscopy signal. *Applied Optics*, 42(16):2915–2921.

Oldfield, R. C. (1971). The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia*, 9(1):97–113.

Olivier, E., Davare, M., Andres, M., and Fadiga, L. (2007). Precision grasping in humans: from motor control to cognition. *Current Opinion in Neurobiology*, 17(6):644–648.

Onose, G., Grozea, C., Anghelescu, A., Daia, C., Sinescu, C., Ciurea, A., Spircu, T., Mirea, A., Andone, I., Spânu, A., et al. (2012). On the feasibility of using motor imagery EEG-based brain–computer interface in chronic tetraplegics for assistive robotic arm control: a clinical test and long-term post-trial follow-up. *Spinal cord*, 50(8):599–608.

Onton, J. and Makeig, S. (2006). Information-based modeling of event-related brain dynamics. *Progress in Brain Research*, 159:99–120.

Oppenheim, A. V. (1999). *Discrete-time signal processing*. Pearson Education, London (UK).

Ortega, P., Colas, C., and Faisal, A. A. (2018). Compact convolutional neural networks for multi-class, personalised, closed-loop EEG-BCI. In *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*, pages 136–141. IEEE.

Ortega, P. and Faisal, A. (2021). HemCNN: Deep Learning enables decoding of fNIRS cortical signals in hand grip motor tasks. In *10th Interational IEEE EMBS Conference on Neural Engineering*. IEEE.

Ortega, P., Zhao, T., and Faisal, A. A. (2020). HYGRIP: Full-stack characterisation of neurobehavioural signals (fNIRS, EEG, EMG, force and breathing) during a bimanual grip force control task. *Frontiers in Neuroscience*, 14:919.

Ostry, D. J. and Feldman, A. G. (2003). A critical evaluation of the force control hypothesis in motor control. *Experimental Brain Research*, 153(3):275–288.

Paek, A. Y., Gailey, A., Parikh, P. J., Santello, M., and Contreras-Vidal, J. L. (2019). Regression-based reconstruction of human grip force trajectories with noninvasive scalp electroencephalography. *Journal of Neural Engineering*, 16(6):66030.

Pan, S. J. and Yang, Q. (2009). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359.

Penfield, W. and Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain*, 60(4):389–443.

Petković, D., Issa, M., Pavlović, N. D., Zentner, L., and Ćojbašić, Ž. (2012). Adaptive neuro fuzzy controller for adaptive compliant robotic gripper. *Expert Systems with Applications*, 39(18):13295–13304.

Pfeifer, M. D., Scholkmann, F., and Labruyère, R. (2018). Signal processing in functional near-infrared spectroscopy (fNIRS): methodological differences lead to different statistical results. *Frontiers in Human Neuroscience*, 11:641.

Pfurtscheller, G., Brunner, C., Schlögl, A., and Da Silva, F. L. (2006). Mu rhythm (de) synchronization and EEG single-trial classification of different motor imagery tasks. *Neuroimage*, 31(1):153–159.

Pfurtscheller, G. and Lopes Da Silva, F. H. (1999). Event-related EEG/MEG synchronization and desynchronization: Basic principles. *Clinical Neurophysiology*, 110(11):1842–1857.

Pfurtscheller, G., Müller, G. R., Pfurtscheller, J., Gerner, H. J., and Rupp, R. (2003). 'Thought'–control of functional electrical stimulation to restore hand grasp in a patient with tetraplegia. *Neuroscience Letters*, 351(1):33–36.

Phan, P., Highton, D., Lai, J., Smith, M., Elwell, C., and Tachtsidis, I. (2016). Multi-channel multi-distance broadband near-infrared spectroscopy system to measure the spatial response of cellular oxygen metabolism and tissue oxygenation. *Biomedical Optics Express*, 7(11):4424–4440.

Phillipou, A., Castle, D. J., Abel, L. A., Gurvich, C., and Rossell, S. L. (2018). An overlooked brain region in the aetiology of anorexia nervosa: The importance of behaviourally driven neuroimaging analysis. *Journal of Experimental Neuroscience*, 12:1179069518820068.

Phillips, C. G. (1977). Corticospinal neurones. Their role in movement. *Quarterly Journal of Experimental Physiology and Cognate Medical Sciences: Translation and Integration*, 63(3):291–292.

Pineiro, R., Pendlebury, S., Johansen-Berg, H., and Matthews, P. (2002). Altered hemodynamic responses in patients after subcortical stroke measured by functional MRI. *Stroke*, 33(1):103–109.

Pinti, P., Aichelburg, C., Gilbert, S., Hamilton, A., Hirsch, J., Burgess, P., and Tachtsidis, I. (2018). A review on the use of wearable functional near-infrared spectroscopy in naturalistic environments. *Japanese Psychological Research*, 60(4):347–373.

Pinti, P., Scholkmann, F., Hamilton, A., Burgess, P., and Tachtsidis, I. (2019). Current status and issues regarding pre-processing of fNIRS neuroimaging data: an investigation of diverse signal filtering methods within a general linear model framework. *Frontiers in Human Neuroscience*, 12:505.

Pistohl, T., Ball, T., Schulze-Bonhage, A., Aertsen, A., and Mehring, C. (2008). Prediction of arm movement trajectories from ECoG-recordings in humans. *Journal of Neuroscience Methods*, 167(1):105–114.

Pistohl, T., Schmidt, T. S. B., Ball, T., Schulze-Bonhage, A., Aertsen, A., and Mehring, C. (2013). Grasp detection from human ECoG during natural reach-to-grasp movements. *PloS One*, 8(1):e54658.

Plichta, M., Heinzel, S., Ehlis, A.-C., Pauli, P., and Fallgatter, A. (2007). Model-based analysis of rapid event-related functional near-infrared spectroscopy (NIRS) data: a parametric validation study. *Neuroimage*, 35(2):625–634.

Plichta, M., Herrmann, M., Baehne, C., Ehlis, A.-C., Richter, M., Pauli, P., and Fallgatter, A. (2006). Event-related functional near-infrared spectroscopy (fNIRS): are the measurements reliable? *Neuroimage*, 31(1):116–124.

Plow, E. B., Cunningham, D. A., Beall, E., Jones, S., Wyant, A., Bonnett, C., Yue, G. H., Lowe, M., Wang, X.-F., Sakaie, K., et al. (2013). Effectiveness and neural mechanisms associated with tDCS delivered to premotor cortex in stroke rehabilitation: study protocol for a randomized controlled trial. *Trials*, 14(1):1–10.

Porter, R. and Lemon, R. (1993). *Corticospinal function and voluntary movement*. Oxford University Press, Oxford (UK).

Rao, S. M., Binder, J., Bandettini, P., Hammeke, T., Yetkin, F., Jesmanowicz, A., Lisk, L., Morris, G., Mueller, W., Estkowski, L., et al. (1993). Functional magnetic resonance imaging of complex human movements. *Neurology*, 43(11):2311–2311.

Ray, S., Crone, N. E., Niebur, E., Franaszczuk, P. J., and Hsiao, S. S. (2008). Neural correlates of high-gamma oscillations (60-200 Hz) in macaque local field potentials and their potential implications in electrocorticography. *Journal of Neuroscience*, 28(45):11526–11536.

Rayegani, S., Raeissadat, S., Sedighipour, L., Mohammad Rezazadeh, I., Bahrami, M., Eliaspour, D., and Khosrawi, S. (2014). Effect of neurofeedback and electromyographic-biofeedback therapy on improving hand function in stroke patients. *Topics in Stroke Rehabilitation*, 21(2):137–151.

Renton, T., Tibbles, A., and Topolovec-Vranic, J. (2017). Neurofeedback as a form of cognitive rehabilitation therapy following stroke: A systematic review. *PloS One*, 12(5):e0177290.

Roland, P. E., Larsen, B., Lassen, N. A., and Skinhoj, E. (1980). Supplementary motor area and other cortical areas in organization of voluntary movements in man. *Journal of Neurophysiology*, 43(1):118–136.

Rosenbaum, D. A. (2009). *Human motor control*. Academic press, San Diego, CA (USA).

Royer, A. S., Doud, A. J., Rose, M. L., and He, B. (2010). EEG control of a virtual helicopter in 3-dimensional space using intelligent control strategies. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 18(6):581–589.

Ryun, S., Kim, J. S., Jeon, E., and Chung, C. K. (2017). Movement-related sensorimotor high-gamma activity mainly represents somatosensory feedback. *Frontiers in Neuroscience*, 11:408.

Sachs, N. A., Ruiz-Torres, R., Perreault, E. J., and Miller, L. E. (2015). Brain-state classification and a dual-state decoder dramatically improve the control of cursor movement through a brain-machine interface. *Journal of Neural Engineering*, 13(1):016009.

Sadaghiani, S. and Kleinschmidt, A. (2016). Brain networks and α-oscillations: structural and functional foundations of cognitive control. *Trends in Cognitive Sciences*, 20(11):805–817.

Sanger, T. D. (2003). Neural population codes. *Current Opinion in Neurobiology*, 13(2):238–249.

Sauseng, P., Gerloff, C., and Hummel, F. C. (2013). Two brakes are better than one: the neural bases of inhibitory control of motor memory traces. *Neuroimage*, 65:52–58.

Sauseng, P., Klimesch, W., Gerloff, C., and Hummel, F. C. (2009). Spontaneous locally restricted EEG alpha activity determines cortical excitability in the motor cortex. *Neuropsychologia*, 47(1):284–288.

Scarpa, F., Cutini, S., Scatturin, P., Dell'Acqua, R., and Sparacino, G. (2010). Bayesian filtering of human brain hemodynamic activity elicited by visual short-term maintenance recorded through functional near-infrared spectroscopy (fNIRS). *Optics Express*, 18(25):26550–26568.

Schalk, G., Kubanek, J., Miller, K., Anderson, N., Leuthardt, E., Ojemann, J., Limbrick, D., Moran, D., Gerhardt, L., and Wolpaw, J. (2007). Decoding two-dimensional movement trajectories using electrocorticographic signals in humans. *Journal of Neural Engineering*, 4(3):264.

Schirrmeister, R. T., Springenberg, J. T., Fiederer, L. D. J., Glasstetter, M., Eggensperger, K., Tangermann, M., Hutter, F., Burgard, W., and Ball, T. (2017). Deep learning with convolutional neural networks for EEG decoding and visualization. *Human Brain Mapping*, 38(11):5391–5420.

Schroeter, M. L., Bücheler, M. M., Müller, K., Uludağ, K., Obrig, H., Lohmann, G., Tittgemeyer, M., Villringer, A., and von Cramon, D. Y. (2004). Towards a standard analysis for functional near-infrared imaging. *Neuroimage*, 21(1):283–290.

Scott, S. H. (2016). A functional taxonomy of bottom-up sensory feedback processing for motor actions. *Trends in Neurosciences*, 39(8):512–526.

Scott, S. H., Brown, I. E., and Loeb, G. E. (1996). Mechanics of feline soleus: I. Effect of fascicle length and velocity on force output. *Journal of Muscle Research & Cell Motility*, 17(2):207–219.

Seghier, M. L. (2008). Laterality index in functional MRI: methodological issues. *Magnetic Resonance Imaging*, 26(5):594–601.

Senior, A. W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., Qin, C., Žídek, A., Nelson, A. W., Bridgland, A., et al. (2020). Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710.

Sensoy, M., Kaplan, L., Cerutti, F., and Saleki, M. (2020). Uncertainty-aware deep classifiers using generative models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5620–5627.

Sensoy, M., Kaplan, L., and Kandemir, M. (2018). Evidential deep learning to quantify classification uncertainty. *arXiv preprint arXiv:1806.01768*.

Serrien, D. J., Ivry, R. B., and Swinnen, S. P. (2006). Dynamics of hemispheric specialization and integration in the context of motor control. *Nature Reviews Neuroscience*, 7(2):160–166.

Shah, S. A., Tan, H., and Brown, P. (2017). Continuous force decoding from deep brain local field potentials for Brain Computer Interfacing. In *2017 8th International IEEE/EMBS Conference on Neural Engineering*, pages 371–374. IEEE.

Shah, S. A., Tan, H., Tinkhauser, G., and Brown, P. (2018). Towards real-time, continuous decoding of gripping force from deep brain local field potentials. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 26(7):1460–1468.

Shibusawa, M., Takeda, T., Nakajima, K., Ishigami, K., and Sakatani, K. (2009). Functional near-infrared spectroscopy study on primary motor and sensory cortex response to clenching. *Neuroscience Letters*, 449(2):98–102.

Shibuya, K., Kuboyama, N., and Yamada, S. (2016). Complementary activation of the ipsilateral primary motor cortex during a sustained handgrip task. *European Journal of Applied Physiology*, 116(1):171–178.

Shibuya, K., Sadamoto, T., Sato, K., Moriyama, M., and Iwadate, M. (2008). Quantification of delayed oxygenation in ipsilateral primary motor cortex compared with contralateral side during a unimanual dominant-hand motor task using near-infrared spectroscopy. *Brain Research*, 1210:142–147.

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359.

Silvoni, S., Ramos-Murguialday, A., Cavinato, M., Volpato, C., Cisotto, G., Turolla, A., Piccione, F., and Birbaumer, N. (2011). Brain-computer interface in stroke: a review of progress. *Clinical EEG and Neuroscience*, 42(4):245–252.

Simonyan, K., Vedaldi, A., and Zisserman, A. (2013). Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*.

Sitaram, R., Zhang, H., Guan, C., Thulasidas, M., Hoshi, Y., Ishikawa, A., Shimizu, K., and Birbaumer, N. (2007). Temporal classification of multichannel near-infrared spectroscopy signals of motor imagery for developing a brain-computer interface. *Neuroimage*, 34(4):1416–1427.

Snoek, G. J., IJzerman, M. J., Hermens, H. J., Maxwell, D., and Biering-Sorensen, F. (2004). Survey of the needs of patients with spinal cord injury: Impact and priority for improvement in hand function in tetraplegics. *Spinal cord*, 42(9):526–532.

Soekadar, S. R., Silvoni, S., Cohen, L. G., and Birbaumer, N. (2015). Brain–machine interfaces in stroke neurorehabilitation. In *Clinical Systems Neuroscience*, pages 3–14. Springer, Tokyo (Japan).

Soekadar, S. R., Witkowski, M., Mellinger, J., Ramos, A., Birbaumer, N., and Cohen, L. G. (2011). ERD-based online brain-machine interfaces (BMI) in the context of neurorehabilitation: Optimizing BMI learning and performance. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(5):542–549.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.

Stilling, J. M., Duszynski, C. C., Oni, I., Paxman, E., Dunn, J. F., and Debert, C. T. (2019). Using functional near-infrared spectroscopy to study the effect of repetitive transcranial magnetic stimulation in concussion: a two-patient case study. *Frontiers in Neurology*, 10:476.

Subhani, A. R., Mumtaz, W., Saad, M. N. B. M., Kamel, N., and Malik, A. S. (2017). Machine learning framework for the detection of mental stress at multiple levels. *IEEE Access*, 5:13545–13556.

Sussillo, D., Nuyujukian, P., Fan, J. M., Kao, J. C., Stavisky, S. D., Ryu, S., and Shenoy, K. (2012). A recurrent neural network for closed-loop intracortical brain–machine interface decoders. *Journal of Neural Engineering*, 9(2):026027.

Suto, T., Ito, M., Uehara, T., Ida, I., Fukuda, M., and Mikuni, M. (2002). Temporal characteristics of cerebral blood volume change in motor and somatosensory cortices revealed by multichannel near-infrared spectroscopy. In *International Congress Series*, volume 1232, pages 383–388. Elsevier.

Tachtsidis, I. and Scholkmann, F. (2016). False positives and false negatives in functional near-infrared spectroscopy: Issues, challenges, and the way forward. *Neurophotonics*, 3(3):031405–031405.

Takeuchi, N. and Izumi, S.-I. (2013). Rehabilitation with poststroke motor recovery: a review with a focus on neural plasticity. *Stroke Research and Treatment*, 2013.

Tang, J., Coleman, J. E., Dai, X., and Jiang, H. (2016). Wearable 3-D photoacoustic tomography for functional brain imaging in behaving rats. *Scientific Reports*, 6(1):1–10.

Thickbroom, G., Phillips, B., Morris, I., Byrnes, M., and Mastaglia, F. (1998). Isometric force-related activity in sensorimotor cortex measured with functional MRI. *Experimental Brain Research*, 121(1):59–64.

Thomik, A. A., Haber, D., and Faisal, A. A. (2013). Real-time movement prediction for improved control of neuroprosthetic devices. In *2013 6th International IEEE/EMBS Conference on Neural Engineering*, pages 625–628. IEEE.

Todorov, E. (2000a). Direct cortical control of muscle activation in voluntary arm movements: A model. *Nature Neuroscience*, 3(4):391–398.

Todorov, E. (2000b). Reply to 'One motor cortex, two different views'. *Nature Neuroscience*, 3(10):963–964.

Tong, Y., Hocke, L. M., and Frederick, B. B. (2019). Low frequency systemic hemodynamic "noise" in resting state BOLD fMRI: characteristics, causes, implications, mitigation strategies, and applications. *Frontiers in Neuroscience*, 13:787.

Torricelli, A., Contini, D., Pifferi, A., Caffini, M., Re, R., Zucchelli, L., and Spinelli, L. (2014). Time domain functional NIRS imaging for human brain mapping. *Neuroimage*, 85:28–50.

Trakoolwilaiwan, T., Behboodi, B., Lee, J., Kim, K., and Choi, J.-W. (2017). Convolutional neural network for high-accuracy functional near-infrared spectroscopy in a brain–computer interface: three-class classification of rest, right-, and left-hand motor execution. *Neurophotonics*, 5(1):011008.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., and Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage*, 15(1):273–289.

Úbeda, A., Azorín, J. M., Farina, D., and Sartori, M. (2018). Estimation of neuromuscular primitives from EEG slow cortical potentials in incomplete spinal cord injury individuals for a new class of brain-machine interfaces. *Frontiers in Computational Neuroscience*, 12:3.

van den Broek, S. P., Reinders, F., Donderwinkel, M., and Peters, M. (1998). Volume conduction effects in EEG and MEG. *Electroencephalography and clinical neurophysiology*, 106(6):522–534.

Vanzetta, I. and Grinvald, A. (2008). Coupling between neuronal activity and microcirculation: Implications for functional brain imaging. *HFSP Journal*, 2(2):79–98.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.

Venkatakrishnan, A., Francisco, G. E., and Contreras-Vidal, J. L. (2014). Applications of brain–machine interface systems in stroke recovery and rehabilitation. *Current Physical Medicine and Rehabilitation Reports*, 2(2):93–105.

Verstynen, T., Diedrichsen, J., Albert, N., Aparicio, P., and Ivry, R. B. (2005). Ipsilateral motor cortex activity during unimanual hand movements relates to task complexity. *Journal of Neurophysiology*, 93(3):1209–1222.

Vidal, J. J. (1973). Toward direct brain-computer communication. *Annual Review of Biophysics and Bioengineering*, 2(1):157–180.

Vidal, J. J. (1977). Real-time detection of brain events in EEG. *Proceedings of the IEEE*, 65(5):633–641.

Villringer, A., Planck, J., Hock, C., Schleinkofer, L., and Dirnagl, U. (1993). Near infrared spectroscopy (NIRS): a new tool to study hemodynamic changes during activation of brain function in human adults. *Neuroscience Letters*, 154(1-2):101–104.

Völker, M., Schirrmeister, R. T., Fiederer, L. D., Burgard, W., and Ball, T. (2018). Deep transfer learning for error decoding from non-invasive EEG. In *2018 6th International Conference on Brain-Computer Interface (BCI)*, pages 1–6. IEEE.

Volkmann, J., Schnitzler, A., Witte, O., and Freund, H.-J. (1998). Handedness and asymmetry of hand representation in human motor cortex. *Journal of Neurophysiology*, 79(4):2149–2154.

Walker, I. (2015). Deep convolutional neural networks for brain computer interface using motor imagery. Master's thesis, Imperial College of Science, Technology and Medicine. Department of computing, London (UK).

Wang, C., Ulbert, I., Schomer, D. L., Marinkovic, K., and Halgren, E. (2005). Responses of human anterior cingulate cortex microdomains to error detection, conflict monitoring, stimulus-response mapping, familiarity, and orienting. *Journal of Neuroscience*, 25(3):604–613.

Wang, K., Wang, Z., Guo, Y., He, F., Qi, H., Xu, M., and Ming, D. (2017). A brain-computer interface driven by imagining different force loads on a single hand: An online feasibility study. *Journal of Neuroengineering and Rehabilitation*, 14(1):93.

Wang, T., Mantini, D., and Gillebert, C. R. (2018). The potential of real-time fMRI neurofeedback for stroke rehabilitation: a systematic review. *Cortex*, 107:148–165.

Wang, X.-J. (2010). Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological Reviews*, 90(3):1195–1268.

Ward, N., Brown, M., Thompson, A., and Frackowiak, R. (2003). Neural correlates of motor recovery after stroke: A longitudinal fMRI study. *Brain*, 126(11):2476–2496.

Wei, X., Ortega, P., and Faisal, A. (2021). Inter-subject Deep Transfer Learning for Motor Imagery EEG Decoding. In *10th Interational IEEE EMBS Conference on Neural Engineering*. IEEE.

Weichwald, S., Meyer, T., Özdenizci, O., Schölkopf, B., Ball, T., and Grosse-Wentrup, M. (2015). Causal interpretation rules for encoding and decoding models in neuroimaging. *Neuroimage*, 110:48–59.

Westling, G. and Johansson, R. (1984). Factors influencing the force control during precision grip. *Experimental Brain Research*, 53(2):277–284.

Wexler, B. E., Fulbright, R. K., Lacadie, C. M., Skudlarski, P., Kelz, M. B., Constable, R. T., and Gore, J. C. (1997). An fMRI study of the human cortical motor system response to increasing functional demands. *Magnetic Resonance Imaging*, 15(4):385–396.

WHO (2002). *The world health report 2002: reducing risks, promoting healthy life*. World Health Organization, Geneva (Switzerland).

WHO (2011). *World report on disability 2011*. World Health Organization, Geneva (Switzerland).

Winstein, C. and Varghese, R. (2018). Been there, done that, so what's next for arm and hand rehabilitation in stroke? *NeuroRehabilitation*, 43(1):3–18.

Wolpaw, J. R., Birbaumer, N., Heetderks, W. J., McFarland, D. J., Peckham, P. H., Schalk, G., Donchin, E., Quatrano, L. A., Robinson, C. J., and Vaughan, T. M. (2000). Brain-computer interface technology: A review of the first international meeting. *IEEE Transactions on Rehabilitation Engineering*, 8(2):164–173.

Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G., and Vaughan, T. M. (2002). Brain-computer interfaces for communication and control. *Clinical Neurophysiology*, 113(6):767–791.

Wolpaw, J. R., Loeb, G. E., Allison, B. Z., Donchin, E., do Nascimento, O. F., Heetderks, W. J., Nijboer, F., Shain, W. G., and Turner, J. N. (2006). BCI meeting 2005-workshop on signals and recording methods. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 14(2):138–141.

Wolpaw, J. R. and McFarland, D. J. (2004). Control of a two-dimensional movement signal by a noninvasive brain-computer interface in humans. *Proceedings of the National Academy of Sciences*, 101(51):17849–17854.

Womelsdorf, T., Johnston, K., Vinck, M., and Everling, S. (2010). Theta-activity in anterior cingulate cortex predicts task rules and their adjustments following errors. *Proceedings of the National Academy of Sciences*, 107(11):5248–5253.

Wray, S., Cope, M., Delpy, D. T., Wyatt, J. S., and Reynolds, E. O. R. (1988). Characterization of the near infrared absorption spectra of cytochrome AA3 and haemoglobin for the non-invasive monitoring of cerebral oxygenation. *Biochimica et Biophysica Acta Bioenergetics*, 933(1):184–192.

Wriessnegger, S. C., Kirchmeyr, D., Bauernfeind, G., and Müller-Putz, G. R. (2017). Force related hemodynamic responses during execution and imagery of a hand grip task: A functional near infrared spectroscopy study. *Brain and Cognition*, 117:108–116.

Wulsin, D. F., Gupta, J. R., Mani, R., Blanco, J. A., and Litt, B. (2011). Modeling electroencephalography waveforms with semi-supervised deep belief nets: fast classification and anomaly measurement. *Journal of Neural Engineering*, 8(3):036015.

Xia, J., Yao, J., and Wang, L. V. (2014). Photoacoustic tomography: principles and advances. *Electromagnetic Waves (Cambridge, MA)*, 147:1.

Xiloyannis, M., Gavriel, C., Thomik, A. A., and Faisal, A. A. (2015). Gaussian process regression for accurate prediction of prosthetic limb movements from the natural kinematics of intact limbs. In *2015 7th International IEEE/EMBS Conference on Neural Engineering*, pages 659–662. IEEE.

Xiloyannis, M., Gavriel, C., Thomik, A. A., and Faisal, A. A. (2017). Gaussian process autoregression for simultaneous proportional multi-modal prosthetic control with natural hand kinematics. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(10):1785–1801.

Xu, J., Ejaz, N., Hertler, B., Branscheidt, M., Widmer, M., Faria, A. V., Harran, M. D., Cortes, J. C., Kim, N., Celnik, P. A., et al. (2017). Separable systems for recovery of finger strength and control after stroke. *Journal of Neurophysiology*, 118(2):1151–1163.

Yazdan-Shahmorad, A., Kipke, D. R., and Lehmkuhle, M. J. (2013). High gamma power in ECoG reflects cortical electrical stimulation effects on unit activity in layers V/VI. *Journal of Neural Engineering*, 10(6):066002.

Ye, J. C., Tak, S., Jang, K. E., Jung, J., and Jang, J. (2009). NIRS-SPM: Statistical parametric mapping for near-infrared spectroscopy. *Neuroimage*, 44(2):428–447.

Yin, X., Xu, B., Jiang, C., Fu, Y., Wang, Z., Li, H., and Shi, G. (2015). A hybrid BCI based on EEG and fNIRS signals improves the performance of decoding motor imagery of both force and speed of hand clenching. *Journal of Neural Engineering*, 12(3):036004.

Yücel, M. A., Selb, J., Aasted, C. M., Lin, P.-Y., Borsook, D., Becerra, L., and Boas, D. A. (2016). Mayer waves reduce the accuracy of estimated hemodynamic response functions in functional near-infrared spectroscopy. *Biomedical Optics Express*, 7(8):3078–3088.

Zafar, A., Ghafoor, U., Yaqub, M., and Hong, K.-S. (2019). Initial-dip-based classification for fNIRS-BCI. In *Neural Imaging and Sensing 2019*, volume 10865, page 108651N. International Society for Optics and Photonics.

Zhang, P., Li, L., Lin, L., Hu, P., Shi, J., He, Y., Zhu, L., Zhou, Y., and Wang, L. V. (2018). High-resolution deep functional imaging of the whole mouse brain by photoacoustic computed tomography in vivo. *Journal of biophotonics*, 11(1):e201700024.

Zhang, Q., Strangman, G. E., and Ganis, G. (2009). Adaptive filtering to reduce global interference in non-invasive NIRS measures of brain activation: How well and when does it work? *Neuroimage*, 45(3):788–794.

Zhou, Z.-H. (2012). *Ensemble methods: foundations and algorithms*. CRC Press, Boca Raton, FL (USA).

Zwillinger, D. and Kokoska, S. (1999). *Standard probability and statistics tables and formulae*. CRC Press, Boca Raton, FL (USA).

Index

Cnnatt vs linear, 147

Alpha band, 55
Attention mechanism, 133, 152

Banana-shaped scattering, 37
Baseline correction, 90
Batch, 97
Behaviourally driven neuroimaging, 162
Beta band, 55
Bimanual, 116, 128
Bivariate Gaussians, 134
Brain-computer interface, 33, 34
Breathing, 76

Central nervous system, 32
Classification, 93
Closed-loop decoder, 35
Confounds, 83
Convolutional filter, 93
Convolutional kernel, 94
Convolutional layer, 94
Convolutional neural networks, 88, 93
Cortical topography, 58

Data readiness, 76
Dataset, 69, 119
Decoding algorithm, 35
Deep brain stimulation, 48
Deep learning, 66, 67, 128
Deep-fusion, 137, 162
Degrees of freedom, 53, 125, 162, 164
Delta band, 55
Discrete cross-entropy, 94
Discrete force characteristics, 127
Drop-out, 97

Edinburgh inventory, 70
EEG-fNIRS holder, 73
Electrocorticography, 48
Electroencephalography, 43, 49
Electromyography, 48
Electrooculography, 76
End-point forces, 56, 57
Epochs, 79
Event-related desynchronisation, 54
Example, 92
Exponential linear unit, 94

Fatigue, 52
Finger pinch, 53
Fixation point, 72
Force decoding models, 133
Force decoding scheme, 133
Force detection performance, 128, 138
Force onset, 54
Force trajectory reconstruction, 128, 138
Force-movement covariation, 56
Frontoparietal, 60
Functional magnetic resonance, 49
Functional near-infrared spectroscopy, 36, 49

General linear model, 63, 91
Generalisation, 61

Haemodynamic response, 40
Haemoglobin, 38
Hand cross-talk, 151
Hand disentanglement, 128, 140, 145
Hand saliency, 143
Hand transference, 164, 167
HbO/HbR functional signature, 110
HemCNN, 87, 160
HemCNN activations, 107
HemCNN architcture, 96
HemCNN hemi-networks, 94
Hemispheric masking, 94
Hemispherical dominance, 90
Hierarchical principal component analysis, 92
High gamma band, 53
High-level force characteristics, 52, 129
Hilbert features, 131
Homunculus, 32
HYGRiiP, 115, 161
HYGRiP, 69, 88, 159
HYGRiP benchmark, 103

Interhemispheric differences, 110
International 10-20 system, 74
Interpretability, 113, 141, 161, 164

Invasive, 48
Isometric, 54, 61

Kalman filter, 63
Kinematics, 45

Lateralisation, 54, 84, 160
Lateralisation index, 90
Leave-one-subject-out testing, 97
Linear discriminant analysis, 63
Local field potentials, 43, 49
Long short-term memory networks, 65, 133

Maximum hand force saliency, 143
Maximum voluntary contraction, 55, 72
Memory depth, 64
Micro-electrodes, 48
Mixture density network, 133
Modified Beer-Lambert law, 38
Motor neuron, 30
Motor system, 30
Multimodal, 66, 67, 115, 127, 140, 163
Muscle coordination, 59
Muscle precision, 60
Muscle selectivity, 58
Muscle-like force characteristics, 52, 129
Musculoskeletal system, 30

Neural networks, 65
Neurocoupling, 39
Neurofeedback, 164
Neuronal desynchronisation, 84
Non-causal decoder, 64
Non-invasive brain-computer interface, 35
Non-linear decoding, 128
Non-linear modulation, 56, 163
Non-stationary response, 67, 81, 87

Ocular artefacts, 43
Offline decoder, 35
Open-loop decoder, 35
Optical density, 79
Optical intensity, 79
Optode, 37, 73
Overfitting, 64

Parietotemporal, 60
Phase changes, 84
Phasic response, 54
Photoacoustics, 169
Photonics, 168
Preactivation, 94
Primary motor cortex, 32
Primary sensorimotor cortex, 57
Principal component analysis, 92

Raw signals, 131
Real-world tasks, 151, 168
Rehabilitation, 164
Relevance propagation, 99
Residual layers, 133, 152
Respiratory chain, 169

Sample, 92
Sensitivity, 141
Signal to noise ratio, 41, 43, 51, 159
Somatosensory feedback, 54, 56
Spikes, 49
Strenuous force, 160
Stride, 94
Support vector machine, 63
Surrogates, 141
Systemic artefacts, 40, 41

Task generalisation, 61
Time sparsity, 166
Time-sample, 92
Time-scales, 82
Tonic response, 54
Transfer learning, 163, 165
Trial, 92

Uncertainty, 166
Unimanual, 72

Vasodilatory response, 40

Wearability, 49
Wiener filter, 63