Cardiff University
School of Psychology

The role of modelling and computer simulations at various levels of brain organisation

Dominik Krzemiński

Supervisor: Dr Jiaxiang Zhang

May 24, 2021

Abstract

Computational modelling and simulations are critical analytical tools in contemporary neuroscience. Models at various levels of abstraction, corresponding to levels of organisation of the brain, attempt to capture different neuronal or cognitive phenomena. This thesis discusses several models and applies them to behavioural and electrophysiological data.

First, we model a voluntary decision process in a task where two available options carry the same probability of a reward for the outcome. Trial-by-trial accumulation rates are modulated by single-trial EEG features. Hierarchical Bayesian parameter estimation shows that the probability of reward is associated with changes in the speed of accumulation of evidence.

Second, we use a pairwise Maximum Entropy Model (pMEM) to quantify irregularities in the MEG resting-state networks between juvenile myoclonic epilepsy (JME) patients and healthy controls. The JME group exhibited on average fewer local minima of the pMEM energy landscape than controls in the fronto-parietal network. Our results show the pMEM to be descriptive, generative model for characterising atypical functional network properties in brain disorders.

Next, we use a hierarchical drift-diffusion model (HDDM) to study the integration of information from multiple sources. We observe a non-perfect integration in the case of the accumulation of both congruent and incongruent evidence. Based on fitting the HDDM parameters, we hypothesise about the neuronal implementation by extending a biologically plausible neuronal mass model of decision making.

Finally, we propose a spiking neuron model that unifies various components of inferential decision-making systems. The model includes populations corresponding

to anatomical regions, e.g. the dorsolateral prefrontal cortex, orbitofrontal cortex, and basal ganglia. It consists of 8000 neurons and realises dedicated cognitive operations such as weighted valuation of inputs, competition between potential actions, and urgency-mediated modulation.

Overall, this work paves the way for closer integration of theoretical models with behavioural and neuroimaging data.

Acknowledgements

I am most grateful to my supervisor Dr Jiaxiang Zhang for his expert advice, guidance and encouragement. I particularly appreciate him always being available and providing a friendly environment to learn. When I began my PhD, I had limited experience in psychology and decision-making, so I would like to thank Jiaxiang for his patience and giving me the opportunity to grow. Jiaxiang's insightful feedback pushed me to sharpen my thinking and brought my work to a higher level in terms of research methods, but also scientific writing. For all this I am extremely grateful, thank you.

I would like to also thank my advisers Prof. Krish D. Singh and Prof. David Marshall for providing invaluable advice and perspective. I want to thank all my collaborators, in particular Lorena Santamaria, Matthias Treder, Dr Naoki Masuda, Prof. Irena Spasic and especially Prof. Alexander Balinsky. I really appreciate our work on various projects and all insightful discussions.

I am also thankful to all members of the Cognition and Computational Brain Lab, and my friends from Cardiff University Brain Research Imaging Centre who were always there to discuss and challenge my ideas. This work has been supported by the Engineering and Physical Sciences Research Council Scholarship.

Next, my gratitude goes to my parents, Adam and Izabela, and my brother, Sebastian, who always believed in me and were there for me in moments of doubt. *Dziękuję.*

Finally, thank you Ewa for believing in me, for your unconditional love and encouragement. This work would not be possible without you.

Contents

Abreviations

ACC	Anterior Cingulate Cortex
DDM	Drift-Diffusion Model
DIC	Deviance Information Criterion
dlPFC	Dorsolateral Prefrontal Cortex
DMN	Default Mode Network
EEG	Electroencephalography
ERP	Event-Related Potential
fMRI	functional Magnetic Resonance Imaging
FPN	Fronto-parietal Network
ICA	Independent Component Analysis
IDM	Inferential Decision-making
JME	Juvenile Myoclonic Epilepsy
LBA	Linear Ballistic Accumulator
LIP	Lateral Intra-Parietal
LMM	Linear Mixed-Model
LOOIC	leave-one-out information criterion
MEG	Magnetoencephalography
mPFC	medial Prefrontal Cortex
MVPA	Multi-Variate Pattern Analysis
NEF	Neural Engineering Framework
NMM	Neural Mass Model
OFC	Orbitofrontal Cortex

PCA	Principal Component Analysis
pMEM	pairwise Maximum Entropy Model
pSMA	Presupplementary Motor Area
rIFC	right Inferior Frontal Cortex
RT	Response Time
SC	Structural Connectivity
SDT	Signal Detection Theory
SMN	Sensorimotor Network
SVM	Support Vector Machine
TVB	The Virtual Brain

List of Figures

List of Tables

Chapter 1

Introduction

What I cannot create,
I do not understand.

Richard Feynman

1.1 The importance of modelling in neuroscience

The human brain consists of almost 86 billion neurons, each with around 10,000 synapses [290]. Highly complex wiring is complemented with a non-homogeneous hierarchical structure and the nonlinear behaviour of the neuronal cells [61, 184]. All this makes the brain one of the most complicated biological systems to study. Despite decades of intensive research of various levels of its organisation (molecular, cell, circuit, and cognitive), the operating principles of the mind and brain remain largely unknown.

One of the most powerful tools in neuroscientific research is computational modelling [87, 228]. Models typically build upon the experimental evidence to simplify the understanding of complex cognitive phenomena [43]. The importance of computational modelling was emphasised in 2011, when the European Commission selected the *Human Brain Project* as one of its 'Flagship Projects' and contributed around €1 billion to 'advance knowledge in the fields of neuroscience, computing, and brain-related medicine' [260]. The initial ambitious goal of this initiative was to create a collection of

hardware, computational algorithms and databases that would enable comprehensive large-scale brain simulations. The project was an extension of the *Blue Brain Project* (BBP), founded by the Brain and Mind Institute of the École Polytechnique Fédérale de Lausanne [259]. This model simulated the activity of an artificial cellular neocortical column consisting of 10,000 cells [177]. The BBP was one of the first attempts to transfer a part of the brain *in silico*.

The above-mentioned models focus on simulating an idle cognitive state without any specific questions or hypothesis. Instead, more often, the modellers focus on an isolated cognitive problem due to overwhelming complexity of the brain's anatomy and the mind. However, in his seminal report, Allen Newell appealed that testing one hypothesis about the brain at a time would not reveal the holistic cognitive algorithms that govern its function [279]. He argued that hypothesis testing must be supported by the more comprehensive models that are able to perform the cognitive tasks. The *Semantic Pointer Architecture: Unified Network* (SPAUN) aims to address this issue. The model consists of 2.3 million spiking neurons organised in hierarchical modules inspired by brain anatomy [381]. The model is able to perform six cognitive tasks: digit recognition, tracing from memory, serial working memory, simple question answering, addition by counting, and symbolic pattern completion [113]. SPAUN is general enough to switch between the tasks without any external 'rewiring'.

Today, scientists emphasise more than ever the key points of Newell's arguments [228, 43, 239]. In 1980s, the cognitive science made huge advances using 'symbolic' and 'connectionist' architectures to implement task-performing cognitive models that closely matched behavioural data [280, 265]. Rapid development of neuroimaging techniques (EEG, MEG, fMRI) enabled measuring of brain activity during active information processing. For example, it is easier now to compare biological plausibility of neuronal models with a detailed evidence from neuroimaging data [51, 347]. Although large-scale brain models allow for more holistic view on the brain activity, specific brain areas or tasks are still not well understood in isolation. For example, basal ganglia is known to play a crucial role in decision-making, but the precise algorithm that rules its behaviour is still not fully discovered [407, 50].

Currently, hundreds of various models exist in neuroscience [228], ranging from de-

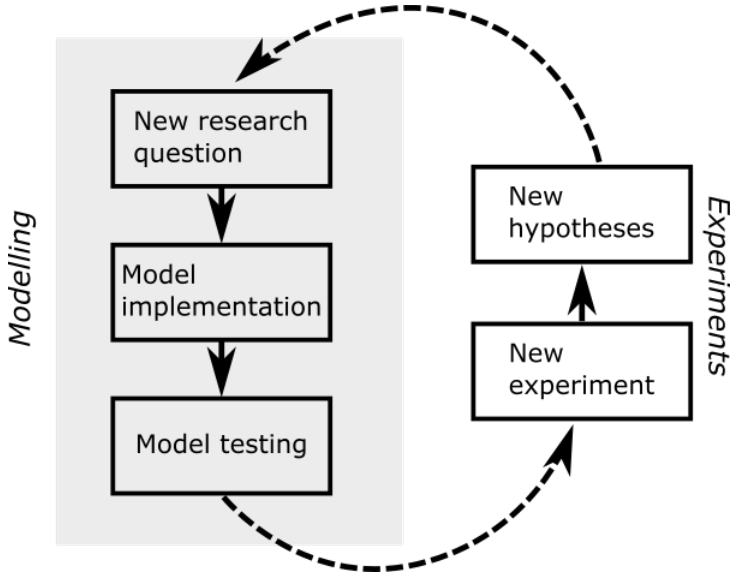

Figure 1.1: *General modelling framework. The model interacts with experimental data through the generation of new model-based predictions, and in turn the experimental data provides new unexplained data to model. The schematic is a simplified version of Figure 1 from [43].*

tailed models of action potential propagation in single neurons [181], through neuronal mass models [454, 140], to large-scale network models of the whole brain [114, 347]. Another branch of models consists of approaches that were only loosely inspired by certain aspects of human cognition or biology of the brain, like high-level models of decision making [326, 49, 470] or artificial neural networks [383].

Despite all those differences, commonalities can be found in the modelling approach. The workflow consists of the following three steps: (1) framing the research question; (2) implementing the model; and (3) testing the model (see Figure 1.1) [43]. All these steps can be further broken down into even more precisely defined components. For example, in [43], a detailed 10-step modelling guide is suggested, including framing the question, formulating specific or theoretic hypotheses, selecting the modelling toolkit, implementing the model, validating it, and more. The most important step, however, is further experimental validation of the existing models. When experimental work brings new hypotheses to a researched problem, the process is then repeated [43, 228]. In this thesis, I aim to apply this pipeline to four different types of models and prove its utility in that context.

The remainder of this chapter is laid out as follows. In Section 1.2, I discuss the definition of a 'model' in the context of neuroscience. Next, I will review selected common taxonomies of cognitive computational models (Sections 1.3 and 1.3.3). Then, in Section 1.4, I will discuss examples of important modelling work at various levels of abstraction. Next, as a few models in this thesis are informed by electrophysiological data, the neural origin of EEG and MEG signals will be described in Section 1.5. At the end (Section 1.6), I will briefly characterise recent developments in a decision-making subfield of neuroscience, as the majority of research presented in this work is centred around this topic.

1.2 Defining 'model'

Before reviewing different aspects of modelling in neuroscience, the proper definition of 'computational model' is needed. *Cambridge Advanced Learner's Dictionary & Thesaurus* describes a 'model' in the mathematical context as 'the activity of using simple mathematical descriptions of a system or process to make calculations or predict what might happen' [313].

This general definition captures the two most important properties of a model: (a) the simplification of the description of a system, and (b) predictive power. For example, statistical tools, like the t-test or ANOVA, can be used for testing hypotheses about the equality of means. For a group-level analysis, this allows answering simple questions about the group difference. Machine learning models fit parameters based on a specific learning rule to solve classification or regression problems [42] (see also Section 1.4.3). This allows predicting a label or a dependent variable from previously unseen data.

On the other hand, we can distinguish *neuronal models* that aim to mimic aspects of the brain's processing of information, at a certain level of abstraction [228, 84]. For instance, in visual neuroscience, deep-learning networks can be a good model of image recognition [73], but poorer for learning representations. This is because during the biological development of an intelligent species, labelled examples are not available for learning in the environment, as they are in supervised machine learning . Another

example includes psychophysical models that map the sensory inputs onto behavioural accuracies or response times, but focus less on on the underlying brain activity [326, 49].

The definition of a 'model' does not explicitly mention the plausibility of the implementation. Indeed, in Section 1.3, I will discuss various computational models at different levels of abstraction, or scale, of the organisation of the brain.

An important aspect of a model is its 'testability'. It must be eligible for verification based on empirical evidence (see the pipeline from Figure 1.1) [228, 43]. A biological realism can be sacrificed to a certain degree in order to generate a model that is easier to understand and verify [239]. As an example, a mechanistic neuronal model does not need to consider all AMPA, GABA and NMDA synaptic connections if it was shown that some of them do not play a critical role in the problem under investigation [454]. Furthermore, the reliability of a computational model is also important. Of course, certain features of such may depend on stochastic parameters, but on average these should return robust predictions. As an example, an option chosen by a cognitive model of decision-making (see Section 1.6 for details) may change on every trial, but estimated values of posterior response time distributions from the whole experiment are consistent between the runs of simulations.

1.3 Overview of modelling approaches

1.3.1 Levels of abstraction

The existing neuronal models can be categorised in a few different ways [250, 228, 261]. The most notable way is a three-level categorisation suggested by Marr [261]. He identified three major stages for modelling a system:

1. The **computational** level: this says what is the overarching strategy behind the system, or simply what is its goal.

2. The **algorithmic** level (also known as the *representational* level): this says how this system might be implemented.

3. The **implementation** level: this says how the system might be precisely implemented using its substrates.

To offer an intuitive understanding of these categories, a simple example of a cash register can be considered [261]. The goal of the cash register needs to be identified first in a top-to-bottom modelling approach. It can be formulated generally as 'quick addition of prices from barcodes of products'. At the representational and algorithmic level, there are a few choices on how to perform such task. First of all, one needs to decide on Roman or Arabic numerals. The two approaches can be tested experimentally. Arabic numerals offer a faster algorithm for addition than their Roman counterparts. For example, Arabic numbers are easily represented as integers of ones, tens, hundreds and so on. Starting from ones, the addition of two numbers can be performed and, if it exceeds 9, the trailing integer is left out and one moves to the tens. The process is repeated for the following numbers in order from right to left. This is not how a real cash register works, but it can be a useful heuristic that helps to understand its efficiency. For the implementation level, one can note that computational devices use electric boards with logic gates for basic operations. For electronic chips, the binary representation of numbers is more convenient than the decimal (i.e. flow or lack of flow of the current). The algorithm of addition remains, however, the same. This simple example demonstrates a typical top–bottom approach to modelling [261, 250].

These three principles were inspired by research into the human visual system [261]. Over the years, these have been generalised to other areas of neuroscience [239, 113, 250]. There are two common approaches to modelling: (a) looking at a problem from the implementation perspective and building upon that (*bottom–up* approach), or (b) finding a high-level description of a problem first and looking for its specific implementations in the brain (*top–down* approach) [43]. For example, cognitive science typically starts from computational theory, decomposing cognition into smaller subsystems (vision, senses, motor control) [271]. As an alternative, certain computational neuroscience approaches may start from the bottom and build blocks of representations and algorithms from biologically plausible computational units [228]. .

Marr's scheme is the most influential categorisation of models, but it has flaws.

The three levels are very general, which indicates its inspiration from engineering, a field that deals with less complicated systems than neuroscience [250]. Recently, various alternatives, not restricted to only three levels, have emerged [239]. Bechtel and Richardon's mechanistic approach can be characterised as a hierarchy of 'levels of mechanism', which does not have a fixed number of stages [29]. For example, a car can be seen as an object consisting of interacting components, such as an engine, steering wheel, brakes, etc. Each component can be further decomposed into its own parts (e.g. the braking system) and so forth. This idea is strongly aligned with the Minsky hypothesis of a 'society of the mind' [271]. This theory models human cognition by building up from the interactions of simple parts, called agents, which are themselves mindless. The model postulates that the interactions of the agents are the fundamental thinking entities from which intelligent behaviour arises [271].

1.3.2 Models of cognition

Another useful taxonomy of the models of cognition is related to their rooting in three influential ideas: 'symbolicism', 'connectionism' and 'dynamicism' [113].

According to the *Philosophy of Mind Dictionary*, the term 'symbolicism' refers to 'an approach to understanding human cognition that is committed to language-like symbolic processing as the best method of explanation.' In that view, thinking is postulated to be done via the manipulation of mental symbols [280]. Formally, symbols are perceptible objects that stand for something else, for example: the alphabet, a road sign, or a red edible fruit (i.e. an apple). This approach enables using the advances of modern engineering and applying them to reverse-engineer the brain [127]. A widely applied cognitive architecture that was born from symbolicism is an Adaptive Control of Thought—Rational (ACT-R) model proposed by Anderson [6]. ACT-R has been used to model memory, attention, and executive control [6, 7, 8]. One of the main challenges of the symbolic view of cognition is its dependence on a number of hand-crafted representations or algorithms underlying certain cognitive processes. In practical application, it requires laboriously pre-defining the logical relationships. In their initial applications, the strategies of symbolicism were advantageous at formal

and language-like tasks, but performed worse at sensory, motor and learning tasks than alternative approaches [164]. Recently, however, methods inspired by 'connectionism' started showing equally good, or better performance on the linguistics tasks [96]. In general, the areas of the brain that rely on procedural, or implicit knowledge (like the motor cortex) are difficult to model within symbolicism framework [113].

'Connectionism' is also known as *Parallel Distributed Processing* (PDP) [265]. PDP models are a class of neurally inspired information processing models that attempt to model information processing the way it takes place in the brain. It suggests that cognitive computations consist of independent processing units connected together serially [265]. Connectionist models have been more successful at building models that express broader generalisibility than symbolic approaches. With the advances in high-performance parallel computational architectures (e.g. Graphical Processing Units [GPU] or Tensor Processing Units [TPU]), this idea has influenced a sub-field of machine learning: *deep learning* [460]. Despite the superficial resemblance of artificial neuronal architectures to biological neurons, the computational principles governing PDP have little in common with the way the brain processes information [460, 229]. To overcome this problem, even the pioneers of 'connectionism' suggest calling PDP models 'brain-style abstract networks' [265].

Another major approach to studying cognitive systems is 'dynamicism', sometimes referred to as an 'embedded' approach to cognition [113]. Dynamicism considers cognition as a non-representational low-dimensional dynamical system [416]. The main argument of dynamicism is that a discretisation of the cognitive system into either symbolic or computational units may lead to its mischaracterisation [112]. Symbolic models are fundamentally dependent on a symbolic representation, and connectionism theories represent concepts via distributed representation in simplified networks. In dynamicism, a cognitive state must be linked with the external environment. Hence, dynamicist modelling describes cognition by emphasising the ongoing real-time interaction with the world [112, 113]. By restricting the descriptions of a cognitive state to a low-dimensional system of differential equations, this approach has created several problems. There are not yet any clear ways of justifying the parameter settings, choosing the equations, or creating the system boundaries within the dynamicism frame-

work [112].

Many of the existing cognitive models are combinations of these approaches [387]. For example, the *Semantic Pointers Architecture* represents symbolic concepts via a vector representation implemented with spiking neurons that respond dynamically to the task [113]. The goal of computational cognitive science is to build a general cognitive system with a detailed understanding of its components. Currently, the high-level cognition models do not offer biologically plausible mechanisms for their implementation [378].

The study of isolated cognitive processes contributed to the greater understanding of certain psychological phenomena. For instance, the Hopfield network provides a good model of associative memory; however, the abstraction of neurons and association rules is far from the realistic biological cell, or plasticity process happening in the brain [80]. The numeric research in theoretical neuroscience has characterised the implementation issues (e.g. information transfer [334], or 'fine-tuning' of neural integrators [225]) of the Hopfield network. Nevertheless, it contributes to a better understanding of how associations are created and what might be a limitation of the algorithmic implementations. This example shows that the focus on the isolated parts of cognition, or neurobiological activity can substantially progress the understanding of the information processing by the brain.

Hence, for the remainder of this thesis, I focus on the isolated cognitive systems (e.g. brain areas accumulating evidence, as in Chapters 2 and 4), or resting-state networks (e.g. default mode network, as in Chapter 3).

1.3.3 Levels of biological realism

This section describes the classification of neuronal models based on the level of biological plausibility. Various representations and their corresponding neural measures (that allow validating these models) will be discussed [250, 239].

Starting from the lowest level of analysis, the ubiquitous model of a single neuron is the Hodgkin–Huxley model [181]. This conductance-based model describes how the action potential in a neuron is propagated along the axon. The model is formulated with

four nonlinear, continuous in time differential equations. The equations approximate the electrical characteristics of excitable neuronal cells [181]. The detailed mathematical formulation of this model was informed by discoveries of the ionic mechanisms, which underly the propagation of the action potential in the squid giant axon [182]. Such conductance-based neuron models reproduce electrophysiological measurements to a high degree of accuracy, but because of their intrinsic complexity they are difficult to analyse at a larger scale [140]. With an increasing number of neurons in a system (e.g. the whole cortical column of a visual cortex), numerical simulations of the spiking activity are too computationally expensive, or too difficult to analyse [259, 177]. Thus, the Hodgkin–Huxley model can be reduced to just two abstract state variables. For instance, the FitzHugh–Nagumo model, despite its simplicity, preserves selected properties of its predecessor [126]. The model can be easily interpreted with a phase-plane analysis by zapping between the branches of the cubic nullcline (for details, see Section 1.4.2).

Phenomenological models consisting of masses of neurons are used for analysing neural spike trains [140, 114]. To make them more numerically efficient, abstractions even simpler than those of the FitzHugh–Nagumo model are used [84]. For example, the *integrate-and-fire* neurons are thresholded models of neuronal firing, derived from equivalent electrical RC circuits. Such RC circuit dynamics represents changes in the voltage of a neuron's soma [140]. The action potential of such a neuron shows no difference between one spike and the next. Hence, its shape, travelling along the axon to a postsynaptic neuron, cannot be used to transmit information [140]. Instead, the action potential represents a binary 'event' (either the spike occurs or it does not occur). The event is fully characterised by the arrival time at the synapse [84]. Therefore, standard integrate-and-fire models make no effort to simulate the actual shape of an action potential. Spikes are treated as contributors to a *firing rate* of the larger neuronal population. More complex adaptive exponential integrate-and-fire models address this problem by introducing an additional activation term with an exponential voltage dependence [193]. Thanks to that, the model is capable of describing standard neuronal firing patterns, e.g. bursting, adapting, or delayed spike initiation.

Building on top of the single-neuron models, the *mean-field* models encompass biophysical features through mathematical derivations from averaged activity of the populations of neurons [140, 404]. Even more abstract mechanistic models, but preserving some level of biological plausibility, are the *neural mass-models*, which are further described in Section 1.4.2. These models are useful when studying the implementation (neuronal) level, but they can also link specific neuronal actions with cognitive phenomena [84, 140].

Finally, *cognitive models* describe human cognitive processes for the purposes of comprehension and prediction [250]. These models do not need to focus on biological plausibility [326, 301] (but they can as shown in Chapter 5, or [113, 33]). However, there are examples of high-level models inspired by neural wiring: for example, the Hopfield network (of recurrent artificial neurons) as a model of an associative memory [42, 80]. A higher-level analysis often involves a loss of information. On the other hand, cognitive models offer explanatory concepts that could not be otherwise formulated at the lower levels [250]. These abstract concepts are important for creating overarching models that inspire more biologically plausible ones. As an example, predictive coding has emerged as an alternative to the traditional concept that the brain integrates information from the outside world through one-directional processes with feature detectors. Instead, predictive coding suggests that the brain is constantly updating hypotheses about the world and predicting sensory information by feedback mechanisms [320]. These two high-level ideas can be instantiated as specific neuronal models and make predictions about neurophysiological experiments.

Recently, a lot of attention has been paid to finding the links between various levels of abstraction in the description of the brain. Yet it is difficult to strike a balance between the model's fidelity to the phenomenon and its neuronal plausibility [239]. For example, one should not criticise an abstract cognitive model for a lack of ion channels, in as much as the Hodgkin–Huxley model cannot account for human visual perception [250]. Different levels of analysis provide different viewpoints on the same phenomena, but often are not equivalent to one another. For instance, sorting as a computational objective can be realised by many sorting algorithms at the algorithmic level. This shows that the many-to-one mapping from higher-level explanations to

lower-level implementations suffers from a reduction or limitation of information [140]. Another example could be *consciousness*, which is called an irreducible entity [403]. As a consequence, studying the neuronal correlates of consciousness is immensely difficult. This may suggest that some phenomena are simply emergent and are impossible, in principle, to explain through lower-level accounts [9].

Overall, finding the bridge between different levels of neuroscientific analysis is one of the most difficult challenges of cognitive science [239, 109]. As an empirical field, neuroscience relies on experimental data, which every year becomes of better quality (increased resolution, senor coverage, etc.) [152]. The influx of new data may provide existing models with additional information, or help to formulate new hypotheses. In this thesis, I present models of isolated brain areas, or cognitive phenomena, with an attempt to make links either between their levels of analysis, or with the neuroimaging data. The high-level model of decision-making used in Chapter 2 was regressed against evidence from EEG recordings. In Chapter 4, we show how the parameters of an abstract perceptual decision model correspond to the dynamics of a biologically plausible neural network.

1.4 The utility of modelling in neuroscience

In Section 1.3.3, I characterised the most common taxonomy of neuronal models. One of the categorisations describes the modelling approaches by the level of biological realism. I started from an overview of detailed models of individual neurons, to higher-level general cognitive models. In this section, I describe several examples of studies that motivated or directly influenced the original research reported in this thesis.

1.4.1 Spiking neuron models

Understanding the behaviour of the brain as a whole starts from understanding the activity of a single neuron. Hence, when making a a biologically plausible simulation of the brain or an area of the brain, it is hard to omit the spiking activity of its building blocks. The brain's nerve cell is electrically excitable and communicates with other

cells via *synapses*. The typical neuron can be divided into three parts: the cell body (also called the *soma*), the dendrites (cellular extensions with multiple branches), and an axon wrapped in a myelin sheath (Figure 1.2) [84]. The function of a typical neuron consists in responding to the signals received by the dendrites, which are then summed in the soma (through an elevated membrane voltage). When a threshold is reached, an action potential is triggered [181]. The information about this 'event' travels down the axon to an axon terminal where it triggers a reaction, which releases a neurotransmitter to the following synapse (in the case of a chemical synapse), or hits the junction (in the case of an electrical synapse). Neurons differ in terms of their shape and size, and can be classified by their morphology and function [236, 136]. Depending on the location and the connections with other structures of the brain, neurons also differ in terms of their excitability. Photosensitive ganglion cells in the retina respond to simple visual stimuli. In the primary visual cortex, simple and complex cells display selective activity for edges, contours, or shades of grey [84]. In the medial temporal lobe, selective neurons respond to more abstract concepts like pictures of faces of celebrities, relatives, or memories [315].

Detailed physiological models of the neuronal cell allow a better understanding of the propagation of the action potential and its impact on the behaviour of the neuron [181]. However, such a level of complexity creates practical numerical problems when a simulation consists of thousands or millions of neurons [114, 347]. Thus, simpler models have been suggested (such as the integrate-and-fire neuron, or the leaky integrate-and-

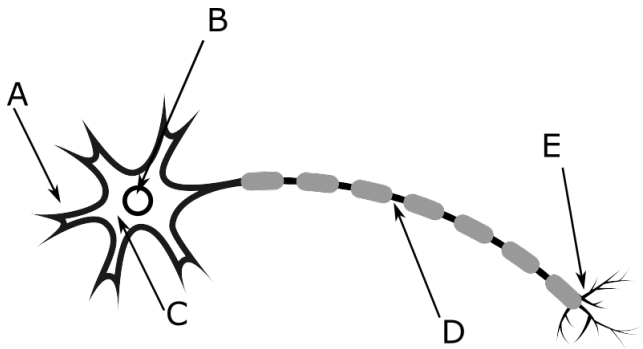

Figure 1.2: *A schematic neuronal cell anatomy: (A) - dendrites, (B) - nucleus, (C) - cell body (soma), (D) - axon, (E) - synaptic terminal.*

fire neuron), that merely approximate the dynamics of a biological neuron [84].

There are several computational frameworks for simulating neural spiking [65, 382]. They differ in their level of biological realism, performance, and number of pre-implemented standard models. One of the most general approaches, the *Neural Engineering Framework* (NEF), was described by Eliasmith and Adnerson in 2003 [114]. NEF does not make any assumptions about what specific function the brain performs, nor about the specifics of its building blocks. Typically, the neuronal ensembles in an NEF model consist of a large number of leaky integrate-and-fire neurons, but these can be customised by the modeller. In the book *How to Build a Brain*, Eliasmith compares the NEF to a Java compiler [113]:

> *"The programmer specifies a program in a high-level language like Java. The Java compiler knows something about the low-level machine language implemented in a given chip, and it translates that high-level description into an appropriate low-level one."*

In this framework, one can specify hypotheses about the function of a specific brain region, and NEF provides a way to connect the neuronal ensembles together in such a way that they realise this function [113]. NEF does not specify *what* the brain computes, but rather *how*. The framework is based on three core principles:

1. **Representation** – neural representations are defined by a combination of non-linear encoding (for example, neuronal tuning curves or spikes) and weighted linear decoding (in space and in time).

2. **Transformation** – neural transformations are functions of variables that are represented by neural populations. Transformations are determined using an alternately weighted linear decoding.

3. **Dynamics** – the dynamics is characterised by considering neural representations as control theoretic state variables. Hence, the dynamics of neurobiological systems can be analysed using control theory.

These three core principles are often complemented by adding noise to the system, something that has been shown to be ubiquitous and a crucial part of neuronal processing [113, 84]. From the above formulation, it is evident that the unit of analysis in NEF is a population (ensemble) of neurons rather than an individual unit.

Since the formulation of the principles of NEF, the framework has been applied to several problems. In [44], an Attentional Routing Circuit model was proposed. It provides a mechanistic description of selective attention processing in the brain. The model simulated three experimental paradigms of an attention mechanism in macaques and displayed a good agreement with the behavioural data and selected aspects of the electrophysiological data [44]. In another study, a spiking model of adaptive control of actions was compared to the activity of the medial prefrontal cortex (mPFC) [33]. In the experimental part of the study, subjects performed a simple task, where the reaction times were improved by learning the expected timing of action-imperative stimuli and preparing movements in advance. The evidence from the neuronal activity of rats' mPFC suggested that neural integration is a key mechanism for adaptive control in precisely timed activities. In [33], it was shown that such a system may consist of coupled neural integrators. The proposed model captured the neural dynamics of the experimental mPFC activity. Moreover, the NEF was used to build the previously mentioned SPAUN model of multiple cognitive phenomena [381].

In Chapter 5, we use NEF to show how to build an anatomically motivated multi-node system of spiking neurons that performs an inferential decision-making task similarly to human agents.

1.4.2 Neuronal mass models

In Section 1.3.2, 'dynamicism' was listed as one of the models of cognition. Dynamicist models are derived from the branch of mathematics called the theory of dynamical systems [84]. Here, I will describe how the theory of dynamical systems can be used to build computational models not only of an individual neuron, but also larger neural populations, that qualitatively exhibits biologically realistic properties. I aim to summarise the key methods and tools from the theory of dynamical systems in the context

of neuroscience [193, 84].

A dynamical system is a system that changes over time [193]. The theory of dynamical systems is the study of the mathematical description of temporally evolving systems [384]. The notion of a system can be as broad as an object falling down under the influence of gravity, or the pattern of the transmission of a disease during a pandemic [277]. In neuroscience, a dynamical system can describe anything from the membrane potential of a neuron to synchronised oscillations in local fields [193].

Dynamical systems can be categorised into *continuous* and *discrete* ones. Continuous systems are formulated by differential equations, while the discrete ones are formulated by iterative maps of difference equations [384]. Continuous nonlinear dynamical systems are the core of computational neuroscience [192, 193], but the discrete ones have applications too, e.g. by means of chaos theory in theoretical research into epilepsy [234].

A general continuous dynamical system θ of dimension k is described by a first order differential equation:

$$\theta = \frac{d\boldsymbol{x}}{dt} = f(\boldsymbol{x}), \tag{1.1}$$

where $\boldsymbol{x} \in \mathbb{R}^k$ and $f : \mathbb{R}^k \to \mathbb{R}^k$. The vector $\boldsymbol{x}$ consists of *dynamical variables*, or *state variables* [193, 384]. For instance, the Hodgkin–Huxley model consists of four dynamical variables that describe a membrane potential and three gating variables for persistent K^+ and transient Na^+ currents [181]. As a result, the evolution of the system is given by a four-dimensional system of ordinary differential equations.

Phase portraits and bifurcation graphs are common tools to analyse a system described by Equation 1.1 [193, 384]. The analysis can be simplified by finding analytical, or numerical solutions for fixed points [384]. Generally, a *fixed point* is an element of the function's domain that is mapped to itself by the function. A fixed point θ^* is asymptotically stable if it attracts all nearby trajectories (in the definition of Poincare) [384]. For instance, for a sufficiently small perturbation near the fixed point $\theta^* + \Delta\theta$, we have $\theta \to \theta^*$ as $t \to \infty$. On the other hand, the fixed point is unstable if any small neighbourhood of the point θ^* contains a trajectory diverging from it. A more rigorous analysis

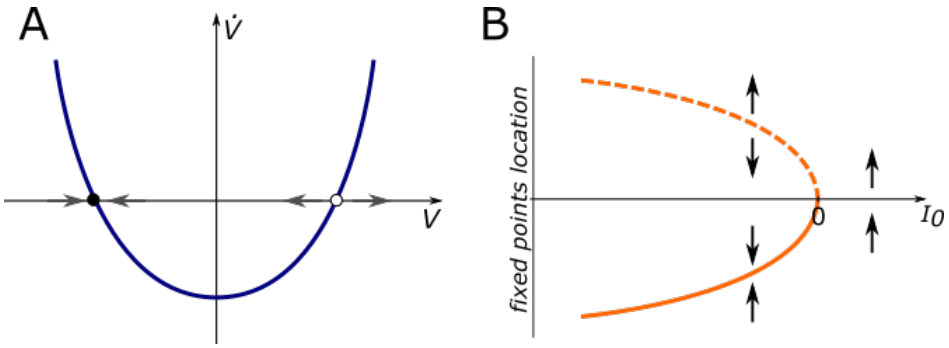

Figure 1.3: *(A) Phase plane for the dynamical system* $\frac{dV}{dt} = V^2 + I_0$ *– the so-called quadratic neuron – with* $I_0 < 0$*. The system has two fixed points: one stable (black circle), one unstable (white circle). The grey arrows show the direction of evolution of* V *in time. (B) The bifurcation diagram for the same system, depending on* I_0*. Dashed line: unstable solution. Solid line: a stable solution.*

of the local stability of fixed points (e.g. considering Lyapunov stability, asymptotic stability, and the Hartman–Grobman Theorem) is discussed elsewhere: [384, 193].

A specific example of the application of the theory of dynamical systems to neuroscience is a simple model of a neuron during quiescence. The model is called a *quadratic neuron* and is a specific instance of the nonlinear integrate-and-fire model [117, 193]. Because this model does not include the propagation of an action potential along the axon, it is an instance of a *point* model. The model serves as a didactic approximation of an exponential integrate-and-fire model depolarised to a state close to repetitive firing. Similar simple point neurons are used for building more complex models with NEF (see Section 1.4.1) [114, 113].

The state variable of a quadratic neuron describes the change of the membrane voltage (V), depending on the external current I_0:

$$\frac{dV}{dt} = V^2 + I_0. \tag{1.2}$$

Note that a solution to this differential equation is the tangent function, which 'explodes' in finite time. For practical applications, a reset value V_r and threshold value V_t are picked, such that when the threshold is reached $V \geq V_t$, the voltage is brought back to $V(t) = V_r$.

For an external current with $I_0 < 0$, the membrane potential V of the neuron evolves towards a fixed value and then does not move away from this value, starting from V_0. This can be seen from a **phase portrait** diagram. An example of such a diagram for the quadratic neuron (a one-dimensional system) has been depicted in Figure 1.3A. The phase portrait contains a graphical representation of a vector field which presents the directions and magnitudes of the changes of a dynamical variable for given coordinates. When $I_0 < 0$, there is a stable steady state at $V^* = -\sqrt{-I_0}$ and an unstable steady state at $V^* = \sqrt{-I_0}$ (see Figure 1.3).

When $I_0 > 0$, the situation is different, as there is no fixed points, meaning that the quadratic neuron will start spiking periodically. The characteristic point where the dynamics of the system changes is at $I_0 = 0$: this is a **bifurcation** point [140]. More formally, a bifurcation describes the appearance of topologically nonequivalent phase portraits under variation of parameters [384]. In other words, a bifurcation occurs when a small smooth change made to the parameter of the system causes a sudden 'qualitative' change in its behaviour (which is reflected in the phase plane). Examples include a sudden change in stability of a steady state or in a rate of oscillation.

For the quadratic neuron, as $I_0 \to 0$ from the negative side, the steady states (V^*) collapse together and at the critical value of $I_0 = 0$ collide to form a single non-hyperbolic steady state. For $I_0 > 0$, there are no real solutions of Equation 1.2, so there are no steady states, see Figure 1.3B.

We distinguish between various types of bifurcations, depending on the types of transitions between their fixed points. The quadratic neuron undergoes a *saddle-node* bifurcation that can generate infinitely slow oscillations [140]. The bifurcation diagram for the quadratic neuron (Figure 1.3B) depicts the number and stability of the systems' solutions depending on a parameter I_0. Other types of bifurcation are also used in neuroscience [384, 84]. For example, an Andronov–Hopf bifurcation models neurons of type II (that emerge with low amplitude but at fixed frequency). In Chapter 4, we analyse the parameters of a dynamical system within the same bifurcation regime, i.e. with the same number and type of fixed points.

The above example of a simple dynamical system is an abstraction of the behaviour of a type I neuron stimulated with an external current in isolation. The brain, how-

ever, contains millions of neurons which are organised in large connected ensembles (called populations or pools) and often exhibit synchronised activity. A mathematical abstraction is made by assuming that the homogeneous populations have the same parameters and receive roughly the same input [84]. The formal assumption of homogeneity is formulated using the following three conditions [140]:

1. all neurons within a pool are identical;
2. all neurons receive the same external input $I_{\text{ext},i}(t) = I_{\text{ext}}(t)$;
3. the interaction strength $w_{i,j}$ the connection between any pair i and j of presynaptic and postsynaptic neurons is 'statistically uniform'.

A method from statistical probability theory helps to reduce such a homogeneous high-dimensional stochastic model to a reduced lower-dimensional model, which approximates the statistical properties of the original one [140]. This approach is called *mean-field reduction* and was first used by Curie and Weiss to study phase transitions [202].

A common neural model that takes advantage of such a simplification is the Wilson–Cowan firing–rate model [449]. In [449], a lower dimensional system was developed which replicated the dynamics of neural populations without directly modelling each neuron. In the context of neuroscience, this approach is called *neural mass modelling* (NMM) [140]. In the Wilson–Cowan model, the averaged firing rates of excitatory and inhibitory pools of neurons are taken into account. Each population is modelled by a single firing rate (a state variable of the dynamical system) bounded between zero (no neurons firing) and one (all neurons in the population activated) [449]. The model assumes that, with no external input, the firing rate of the populations decays to zero, and that the input to the populations is transformed to a firing rate via a sigmoid function.

In a similar approach, the 7200 leaky integrate-and-fire neurons of the lateral intraparietal area (LIP) of the brain were reduced to only two state variables [454]. In [433] a biologically plausible model was proposed to explain – via a series of a recurrent connections mediated by a pool of interneurons – the single-unit recordings and

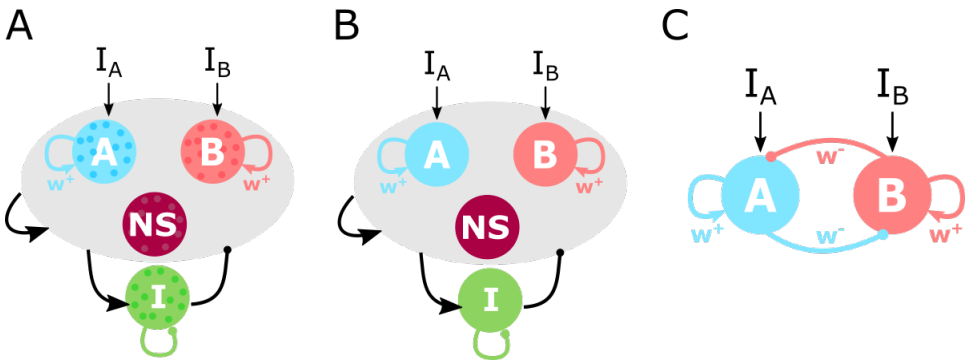

Figure 1.4: *Reduction of a biophysical neuronal decision-making model from [454]. (A) the original model consisting of 7200 leaky integrate-and-fire neurons, endowed with strong recurrent excitation between units and shared inhibition. (B) mean-field approach reduction of 2000 spiking neurons to four neural units (with a total of 11 dynamical variables). (C) final reduced two-variable model with two neural units, endowed with self-excitation and effective mutual inhibition. NS: nonselective excitatory ensembles of cells. I: inhibitory cells. Lines with arrow-heads: excitatory connections. Lines with dot-heads: inhibitory connections.*

behavioural data from a perceptual decision-making experiment (Figure 1.4A). The model agreed well with the experimental data, but, due to its complexity, was difficult to analyse. In [454], it was first assumed that the net input to a neuron in a large homogeneous population can be treated as a Gaussian random process. Following [59], the synaptic currents were approximated by a constant, even though population firing rates typically depend on input currents related to the firing rates. Next, the input–output relation between the interneurons was linearised, based on the typical values of the mean firing rates of inhibitory neurons (8–15 Hz). This allowed reducing the system to 8 state variables with two neuronal pools accumulating evidence over time via self-excitatory connections and recurrent inhibitory activity via interneurons. Finally, it was noted that among the AMPA, NMDA and GABA(A) synapses, the NMDA ones have the longest time constant (on the order of 100 ms) [454]. It was further assumed that all other variables achieve steady states faster than NMDA. Since AMPA and GABA reach steady states much faster, their gating variable becomes proportional to the average firing rates of the presynaptic cells. This allowed replacing the interneuron pool with mutual cross-inhibitory activity coming from the two populations of neurons [454] (see Figure 1.4).

Similar models can be used not only to analyse isolated brain areas, but also to inform large-scale brain simulations. In models of the whole brain, anatomically or functionally defined regions of the brain are represented by a neural mass model, and the state variables of each node are coupled via a connectome to study brain-wide synchronization. The connectivity matrix is derived from structural imaging data (often referred to as *structural connectivity*, SC) [206, 349] or estimated based on functional data using parameter optimisation techniques (so-called *effective connectivity*; EC) [349, 132]. An example of the utility of large-scale brain modelling may be exemplified with the research studying the emergence of functional connectivity resembling empirical data from dynamically shaped slow linear fluctuations [89]. In that study, the even further simplified Wong & Wang model was used – with only a single state variable [454]. Bifurcation diagrams depending on the coupling strength, G, between the nodes of the simulation (representing anatomical regions of the brain) for such a reduced model revealed two bifurcation points: transferring from a single stable firing rate regime to a multi-stable one with a spontaneous state and then to multi-stability with a high firing rate. More importantly though, the correlation between the simulated functional connectivity and the empirical one was maximal for values of G close to the second bifurcation. This was confirmed by the various noise levels of the signal, but also by a simulation with spiking neurons. Overall, the study suggested that FC arises from noise propagation and dynamical slowing down of the fluctuations in an anatomically constrained dynamical system.

Recently, a number of similar studies have been carried out [161, 377] due to the increased availability of software that enables incorporating neuroimaging data with the numerical differential equation simulations for dynamical systems. For example, *The Virtual Brain* (TVB) offers a platform for using SC matrices as the vertices of a graph representing a model of the whole brain. TVB as a simulation environment enables the model-based inference of neurophysiological mechanisms across brain scales [347]. Due to its use of predefined transfer functions, the platform may be used to generate macroscopic neuroimaging signals such as fMRI, EEG or MEG. This allows simulating the averaged brain responses based on the summary statistics of the population, as well as closer personalisation of the models to the individual brain anatomy [347]. It

facilitates an exploration of the consequences of pathological changes in the system, much needed for instance in clinical research [199].

1.4.3 Machine learning

Another family of models commonly used in neuroscience comes from the subfield of artificial intelligence called machine learning [428, 423, 228]. These days, neuroscientists generate enormous amounts of data, matching the complexity of real-world tasks. It has become necessary to find efficient ways to process and analyse these datasets [423].

Machine learning models are non-mechanistic, mainly statistical models [456]. More broadly, these approaches utilise statistics to predict trends and patterns, typically by learning from experience provided in the form of data [423]. Thus, statistical models (including machine learning) are bound to their calibration range and can only predict results within the data space they are calibrated from. Since they are based on correlation and not causality, such models provide limited understanding of the underlying process [456]. This is opposite to the mechanistic models, which are based on the fundamental laws of natural sciences. In these models the physical and biochemical principles constitute the mathematical model description. An essential benefit of mechanistic models is that their parameters have actual physical meaning.

Nevertheless, machine learning helps with multivariate pattern separation (between task conditions or subject groups), which is a common objective of many experimental studies [143], or dimensionality reduction, which is of particular interest as neuroimaging data is usually high-dimensional (and often multi-modal) [423]. Although statistical models do not attempt to directly stereotype the brain, they are useful methods for the analysis of neuroimaging data (as shown in Chapters 2 and 3).

The term *machine learning* was coined by Samuel in 1959, who described it as a 'programming of a digital computer to behave in a way which, if done by human beings or animals, would be described as involving the process of learning' [345]. There are three types of learning procedures: supervised learning, unsupervised learning, and reinforcement learning [428]. The term *learning* in this context means an efficient fitting of the parameters, rather than the cognitive learning process.

Supervised learning

Supervised learning involves a training procedure with access to 'labeled' data. Suppose that we want to build a classifier of animals' images. This means that each picture must have a class assigned to it, for instance, a cat or dog [42]. An example closer to neuroscience could be neuroimaging recordings from a subject watching multiple pictures of animals, where every image (and its associated signal) has a label: a subject looking at a dog, or cat.

For a more specific example, in [358], a supervised classifier, in fact, a support vector machine (SVM), was used to classify the colour of the stimuli watched by subjects from local fMRI brain recordings. The classification was based on the activity of multiple voxels* within the mPFC. A procedure that involves a large number of a signal features (properties that affect the classification process) is often called *multivariate pattern analysis* (MVPA) [143]. Examining the accuracy of the SVM classifier over time and between different brain regions allowed inferring where and when the colour was represented in the brain [358].

Similar approaches are used to detect the intention of a subject during exposure to repeated flashing stimuli on a screen [344]. This type of device is known in brain–computer interface (BCI) research as P300-spellers [232]. A speller consists of an EEG can with electrodes connected to a digital amplifier, which records the brain's activity in real time, and transmits the data to a computer. Based on the signal detected from the scalp voltage (for example, event-related potentials, ERP), the intention of a choice can be inferred and translated to an activity on the computer screen, e.g. writing a letter. Typically, such spellers involve intensive training of a machine learning algorithm that recognises patterns of brain activity for each subject separately [344, 232].

The field of machine learning has been revolutionised by the emergence of deep neural networks, which gave rise to a new subfield: deep learning [460, 229, 353]. Deep neural networks (DNNs) are built with artificial neurons, which are organised in layers, similarly to the cortical organisation of the visual cortex [214]. In analogy to a typical biological neuron, an artificial neuronal unit involves a summation of informa-

*3-dimensional pixels.

tion (transferred from presynpatic to postsynaptic cells) with appropriate weights (the strengths of the synaptic connections): all being thresholded by a non-linear function (imitating the release of an action potential after reaching a certain threshold voltage) [460]. Despite its biological inspirations, artificial neural networks have little in common with realistic neuronal models. Simplifying assumptions have been made to make the learning procedures more numerically efficient. For example, it was suggested that the weights between the layers of the network can be adjusted via a *backpropagation procedure* [460], which computes the gradient of a predefined loss function with respect to the weights for each training example. The weights are further adjusted in the direction opposite to the gradient scaled by a constant (*learning rate*) [42]. Although numerically efficient, backpropagation suffers from a series of problems. The gradient descent procedure is not guaranteed to find a global minimum of the error function, but only a local minimum. Moreover, such a learning process is biologically unrealistic, as in the brain the information travels only in one direction through each neuron (from dendrite to axon). However, the overview [442] offers a more thorough discussion of alternative biologically plausible learning mechanisms that can be approximated by backpropagation.

Despite these limitations, deep learning has become crucial in the advancement of neuroscientific research [228, 165]. Deep neural networks can help with processing large amounts of data. For example, convolutional neural networks 'learn' from training data by passing its features through a multilayered network of simple modules. Such networks progressively abstract the data and extract features, which can be applied to analyse new data. In [214] a deep neural network that classifies images with high accuracy was used. The levels of processing of the network were compared with the activity of various stages of the human visual system. It was shown that the recurrent connections in the network offer better resemblance to the biological system than a similar but feed-forward architecture [214]. In another example, [349], deep neural networks were used to study the structure to function mapping of the human brain. The researchers asked the question whether the modest coupling between SC and FC is a fundamental property of a nervous system or merely a limitation of current brain network models [349]. A five-layer architecture with dropout connections (randomly

removed weights to improve generalisation) was used to predict the FC matrix from the same dimensional SC. The results pointed towards the conclusion that the structure–function coupling in human brain networks is substantially tighter than previously suggested. For instance, brain function predicted from an individual's SC explained significant inter-individual variation in cognitive performance [349].

Regression models are another subcategory of supervised machine learning. In a regression problem, the task is to predict a continuous value based on predictors. For example, the simplest scenario is to predict one output variable by fitting a linear function to the data [42]. Regression models are ubiquitous in neuroscience. For example, they help with decoding the spiking activity of neurons in animal studies [174], or comparing experimental conditions in fMRI studies [312].

Unsupervised learning

Supervised learning often relies on the laborious work of human (or semi-automatic) annotators [370]. For supervised deep learning networks, the order of the size of the dataset must be around thousands of examples to achieve a satisfactory training accuracy (but see also the recent advances in few-shot, or one-shot networks discussed elsewhere: [332]). Obtaining that amount of data is not always possible due to practical limitations, such as, for example, the fatigue of human participants, or expensive neuroimaging procedures (like fMRI). Thus, there is another family of machine learning methods, helping to analyse 'unlabeled' data. The goal of unsupervised learning is to discover hidden structures in the data [42]. Usually, this is achieved by exploring a similarity or redundancy between multivariate patterns. In neuroscience, for example, learning a relevant encoding of stimuli by a neural system may help to represent the parameters of an experiment. Unsupervised machine learning approaches can be used to detect extracellular spikes in that context [174].

A ubiquitous unsupervised method in neuroscience is Principal Component Analysis (PCA), a statistical technique to reduce the dimensionality of a problem, which exploits the redundancy in the data, using second-order statistics [42]. A PCA algorithm creates a new representation of a dataset by means of latent variables (or principal

components) that account for the largest variance of the original data. This allows substantially reducing the dimensionality of the data (often to the first two or three most important components) [42, 201]. Another unsupervised machine learning algorithm is Independent Component Analysis (ICA) [188, 34]. ICA is a common M/EEG data pre-processing technique that allows removing muscle, eye blinks, or heart rhythm artefacts [201]. This can be done by searching for the maximally non-Gaussian combination of the components, assuming that the observed data is generated via a linear sum of independent contributors [188]. Currently, ICA has become one of the default steps in M/EEG data cleaning pipelines [201, 162]. Both the PCA and ICA algorithms have been used in the research for this thesis as pre-processing methods (Chapters 2 and 3).

In the context of neuroimaging, data analysis typically uses supervised and unsupervised methods. Reinforcement learning in another machine learning procedure that involves learning an optimal behaviour of an agent by a set of actions that could be taken in response to a state of the system [428]. It has been heavily inspired by behavioural psychology [367]. Reinforcement learning differs from supervised learning, because no precise mapping between the training examples and the labels is required. Instead, the focus is on finding a balance between exploration and exploitation of the agent [242].

A major difference when comparing machine learning models with previous neuronal modelling approaches is that they explicitly rely on the data. Additionally, most machine learning algorithms are not specific to one type of data type. For example, one can use SVM to predict the ERPs in the BCI setting, but also to discriminate between the activity of brain while subjects are looking at faces and houses based on fMRI features [428, 232]. The assumption-free character of machine learning models can be beneficial, but sometimes problematic. For instance, the mechanistic Wong & Wang model, presented in Section 1.4.2, relies on specific assumptions about the neuronal process that it aims to explain (i.e. perceptual decision making) [454]. Such a model offers not only explanatory power, but also insight into the mechanism of the cognitive process (see also the distinction between mechanistic and statistical models at the beginning of Section 1.4.3). The common criticism of machine learning applica-

tions to neuroscience is the inability to provide an understanding of the process under investigation [428]. Nevertheless, when the data gets sufficiently large and complex, the applications of data mining and machine learning become essential.

Machine learning models interpretability

In recent years, the interpretability and explainability of machine learning models has become a topic of particular interest [101, 168, 183]. It is easier to interpret simple models (e.g. linear regression), whereas more complex models typically yield better performance (e.g. deep convolutional neural networks) [228, 385]. Thus, finding an appropriate trade-off is not a trivial task.

However, even the interpretation of linear classifiers has its pitfalls [101]. In a simple scenario, we have only two observation channels: one that contains the source of information with noise, $x_1(t) = s(t) + \eta(t)$, but the second with just noise $x_2(t) = \eta(t)$. A linear combination of channels with weights can be used to extract the signal $s(t)$ from the data $\mathbf{x}$. For example, the signal is recovered by taking the difference $x_1(t) - x_2(t) = \mathbf{w}^T\mathbf{x}$, where $\mathbf{w}^T = [1, -1]$. Such a classifier gives equal weight to both channels. Therefore, interpreting those weights as evidence that the signal-of-interest $s(t)$ is present in a channel would lead to the wrong conclusions [101]. This issue has been highlighted in [168]. The solution consists of multiplying the weights of a linear classifier by the covariance matrix of the input data. The application of the interpretation of a linear classifier will be presented in Chapter 2 (Figure 2.6).

Recently, great progress has been made towards explaining the predictions of more complex models, such as deep learning networks. For example, the Layer Relevance Propagation (LRP) method was used to interpret the architecture used for the single-trial EEG classification problem [385]. The robustness of these approaches is often doubted. In [102], a didactic example was constructed by attaching a small image of a panda to the fMRI scan of the brain in only one of two available classes. Two methods were tested: LRP and relevance maps [333]. In most cases, both methods ignored the 'panda' feature as a predictor of the class. This and other studies only strengthen concerns about the reliability of the more complex models and attempts at

their interpretation [102, 168]. The need for explainable deep learning in, for example, clinical applications, where the results may inform treatment and triage, cannot be neglected. Over the years, brain research has been an inspiration for the development of algorithmic-level implementations of the general mechanisms of human and animal learning and pushing machine learning researchers towards relevant solutions [165, 428, 143], whereas machine learning algorithms have facilitated the analysis of large amounts of neuronal data [165, 143].

1.4.4 Models informed by neuronal data

So far, mainly theoretical models of neuronal activity (e.g. Hodgkin–Huxeley, FitzHugh-Nagumo) and models that help with the analysis of neuroimaging data (e.g. SVM, PCA) have been discussed. Recently, more studies have aimed to unify the experimental evidence and *a priori* models [239, 228]. As an example, NMMs have been used to determine large-scale network properties via the simulation of local neural populations integrated with long-distance functional and/or structural connectives [347, 89]. Additionally, recent years have brought about a wider use of network theory in applications to neuroscience [25, 270]. The field is evolving rapidly due to computational developments that enable the study of connectivity in fine anatomical detail and the simultaneous interactions between multiple regions [270]. Network theory provides an intuitive framework for studying relationships between interconnected brain areas and helps study their relevance to behaviour [25]. This approach bridges fine-scale and coarse-grained information drawn from real systems to create network models as data representations, which in turn can be used to inform theoretical work.

Chapter 3 will present a phenomenological model that combines network theory with evidence from neuronal data, the *pairwise Maximum-Entropy Model* (pMEM). The pMEM is rooted in statistical physics (where it is known as the Ising model) [60]. In its original formulation, the model describes the behaviour of a magnetic probe in an external magnetic field under different temperature conditions [60, 389]. It consists of a lattice with binary values -1 or $+1$ representing the spin orientation of an atom. This system undergoes a bifurcation at a critical parameter (originally, the temperature)

where the probe loses its magnetic properties due to the random orientation of spins, which cancels the average magnetic moment of the lattice [60]. Traditionally, the model was used in neuroscience to describe the activity of spiking neurons. The maximum entropy approach is a simple abstraction that allows predicting, on average, up to 90% of the available spatial correlation structure of the neurons (for a distribution of states in an ensemble of up to 10 neurons) [390, 400].

Further studies extended the applications of pMEM to fMRI signals [389, 436, 437] or MEG [233]. For example, [436] brought the first evidence that a pairwise interaction parameter of pMEM may accurately describe the resting-state of human brain networks. They applied the model to selected regions of interest (ROIs), as the fitting procedure is numerically intensive. A contemporary standard computational devices allow fitting the model for up to $N = 12$ nodes (although for larger N, approximate solutions are available). The blood-oxygen-level dependent (BOLD) response was thresholded to get a binary activity (similarly to the Ising model) [436]. The model was fitted to two resting state brain networks (i.e. networks active in the presence of no particular task) with high accuracy and robustness. With that it was possible to create a map of the most common brain states (immediate binary activity distributed across the regions; for more information see Chapter 3.2.6), called an *energy landscape* [389]. Furthermore, the functional interaction matrices derived from the pMEM were similar to the anatomical connectivity matrices. These results were the first evidence of the applicability of pMEM to macroscopic brain signals. They extended previous findings that described activity patterns observed in a slice or small parts of the brain [436]. The method was further used to study resting state brain networks during human bistable perception [437], age-related impairment of brain connectivity [120], and autism related differences in node transitions [435].

The pMEM accounts for the spatial correlations in relatively small ensembles of ROIs, but brains respond in both spatial and temporal patterns. Introducing a temporal dependency for only up to the two last steps increases the complexity of the pMEM model [390]. For example, for a system with $N = 4$ nodes, the dimensionality of the spatial problem is $2^N = 16$, whereas for its temporal extension it becomes $2^{3N} = 4,096$. Another issue with the pMEM model, is that for large N, exponentially

more data are needed to accurately estimate the model parameters. Nevertheless, the pMEM might be a useful tool to examine changes in small resting-state brain networks [437, 437, 233].

The pMEM (further discussed in Chapter 3) is only one example of how networks help to combine neuroimaging data with models. The other examples include informing multivariate NMM with empirically derived SC [89] (for more, see Section 1.4.2), studies of 'small worldness' [270], and Dynamic Causal Modelling [132]. However, network approaches are not the only way to combine the neuronal data with models of the brain, or cognition. Another possibility is incorporating the properties of the neuronal data as a regressor to high-level models of behaviour [67, 288, 465].

For example, the model of decision-making, Drift Diffusion Model (DDM) (for more, see Section 1.6), has been embraced by cognitive neuroscience as it accounts for a variety of reaction-time experiments [49]. In [67], EEG recordings provided a single-trial measure of brain activity, i.e., *theta*-band magnitude from frontal electrodes. Instead of estimating one static threshold parameter of DDM per subject across trials, the proposed model assumed that the threshold varied for each trial according to a linear model (as a function of the measured *theta* activity). It was found that this activity was correlated with an increased decision threshold in the selected category of trials. The parameters differed when subjects were undergoing a deep brain stimulation of the subthalamic nucleus. As hypothesised by the model, the subthalamic nucleus was responsible for the increased decision threshold modulated by the increase of the cortical *theta* activity.

Moreover, in [288], single-trial EEG attention measures informed DDM fitting in order to explain the evidence accumulation and decision preparation in a visual decision-making task. A singular value decomposition, similar to PCA, was used to find per-trial estimates of the evoked responses in the EEG. The method boosts the signal to noise ratio in the estimates of the synchronised activity of cortical neurons, but also reduces the multivariate signal (coming from multiple electrodes) to a single latent variable. Next, a hierarchical Bayesian version of DDM was used with a measure of attention as regressor, which consisted of the magnitudes of the single-trial ERP components P200 and N200. This modelling approach shows how evidence from electrophysiology can

be used to understand the influence of visual attention on the mechanisms of decision making [288].

Similar *neurocognitive* models have been used to combine mathematical models of cognition with observations of brain behaviour to explain and predict perceptual decision making [301, 288, 297]. The usefulness of combining behavioural models and neural dynamics has been motivated on theoretical grounds. Furthermore, independent neural measures can significantly improve the explanation and prediction of the cognitive process [301].

1.5 Neuroimaging

Modelling is often informed or tested with recordings of neuronal activity [239, 43, 288]. Depending on the model abstraction, or level of plausibility, various signal modalities can be used (see Section 1.3.3). In this thesis, an emphasis was put on the models of human brain activity. In contemporary neuroimaging, there are a number of non-invasive tools to study the function of the human brain [37, 68, 308, 445, 267]. This section summarises common non-invasive neuroimaging techniques that were employed for the purpose of this research: MRI (in Chapter 3) and fMRI, EEG (in Chapter 2), and MEG (in Chapter 3).

1.5.1 Magnetic Resonance Imaging

One of the most popular neuroimaging tools is the Magnetic Resonance Imaging (MRI) scanner. MRI devices use strong magnetic fields and radio waves to generate images of the organs in the body [200]. Structural MRI is used to uncover anatomical structures, e.g. neurological and other soft tissues. MRI takes advantage of the different factors within the body such as the presence of different chemical bonds, paramagnetic ions and rate of flow of fluids that produce different magnetic resonance signals and hence show contrast between tissues. Functional Magnetic Resonance Imaging (fMRI) focuses on the processes in the human brain as they develop in time. This allows mapping a cognitive function onto specific brain regions in time [37, 200]. For the past 30

years, fMRI has dominated human brain mapping research, because it does not require subjects to undergo any injections, surgery, or to be exposed to ionising radiation [200]. fMRI measures BOLD contrast, which depends on the blood flow (the so-called hemodynamic response function, HRF) related to the energy consumed by a brain area engaged in a specific function [37]. Different magnetic properties of oxygenated or deoxygenated blood affect the relaxation times of the spins aligned via an external magnetic field in the scanner. Despite the good spatial resolution (the standard method can detect changes with a precision of millimeters [308]), fMRI suffers from a poor temporal resolution [200]. Typically, the scans allow sampling the whole brain in time scales of from a second to a few seconds. The human brain, however, performs actions incredibly efficiently. Very simple, perceptual decisions take no more than hundreds of milliseconds to process and execute [470, 288].

1.5.2 Electroencephalography

In many experiments, the exact timing of a task and its neuronal synchronisation cannot be ignored [465, 289]. Recordings of the electrical activity of neurons are much more precise. For many years, researchers have used recordings of individual neurons using clamping techniques, or neuronal ensembles using local field potential (LFP) measures [84, 363]. These usually involved invasive interaction with the brain and surrounding tissue. Unless in a clinical setting, similar techniques are used mainly in animal studies.

Electroencephalography (EEG) provides a fast-sampled non-invasive technique for recording the electrical activity of neurons outside the scalp. The first EEG measurements were performed by Berger [35], and Beck and Cybulski [78]. EEG records the changes of electrical potential difference from the synchronous activity of large population of cortical pyramidal neurons [290, 219]. These neurons differ from the standard neuronal cell. The shape of a pyramidal cell is characterised by a singe large *apical dendrite* that emerges from the apex of a cell [290] (compare with Figure 1.2). The synaptic activity of these neurons generates subtle electrical impulses called *postsynaptic potentials* [163]. EEG waveforms recorded on the scalp are formed due to a

linear superposition of contributions from billions of such micro-current sources. Each source can be approximated by an electrical dipole with a sink around the cell body and a source near the apical dendrite. The primary current (post-synaptic potential) generates a secondary extracellular dipole current. Only synchronised activity (spikes at the same time) of large populations of micro-dipoles can be reliably detected. To record them from a distance, the dipoles must be oriented in the same way, so that the positive and negative charges do not cancel out in the summation [290]. The tissue in the brain is immersed in an ionised *cerebrospinal fluid*. This liquid allows the electromagnetic field of the dipoles to propagate to the surface of the scalp. When recorded from outside the scalp, the electrical signal is attenuated by the skull, skin, and hair [163]. Given the distribution of dipoles within the brain, estimating the external voltage outside the skull is referred to as the *forward problem*, Figure 1.5. The *inverse problem*, in contrast, aims to use the recorded scalp potential distribution to determine the sources' activity and location [290]. In the following formulation, the problem is ill-posed due to the lack of a unique solution. However, with anatomical constraints (coming typically from structural MRI or a priori anatomical knowledge), approximate numerical solutions are possible, for example, Beamforming, or Low Resolution Electromagnetic Tomography [178, 304].

Scalp EEG activity is characterised by oscillations of various frequencies. The frequencies are analysed by means of the power spectral distribution [290]. Several of these oscillations have specific frequency ranges and spatial distributions. Various frequency bands are also associated with the different states of a functioning brain. For example, alpha activity (8–13 Hz) is typically linked with attention and fatigue state [35]. The neuronal networks underlying some of these oscillations are well understood (e.g. the thalamocortical resonance underlying sleep spindles [45], or beta waves [146]).

1.5.3 Magnetoencephalography

According to Maxwell's laws, every electric current is accompanied by a magnetic field. Hence, the current dipoles from the cortical pyramidal neurons generate small magnetic fields [163, 162]. *Magnetoencephalography* (MEG) is a technique that allows

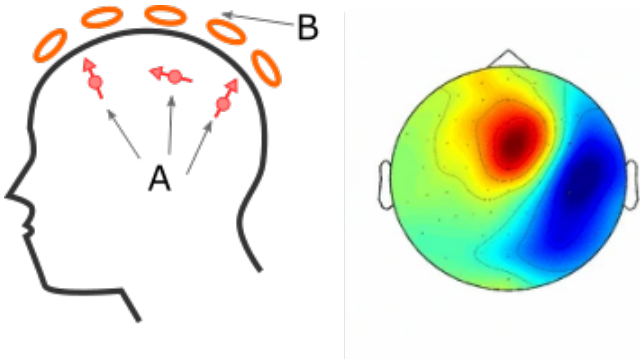

Figure 1.5: *LHS: A schematic of source equivalent dipoles (A) generating M/EEG signals. Multiple dipole generators are captured by the magnetic sensors, or electrodes (B). RHS: Resulting distribution of electric potential or magnetic field on the scalp from a single dipole.*

detecting these fields. The principle of neuronal magnetic signal generation is similar to EEG, but the magnetic field resulting from the intracellular current flow inside the neuronal axon is much greater than the contribution of transmembrane current and passive-current flows [162]. Electrical currents generated by the dipole sources in the cortex produce local changes in the magnetic field outside the scalp. The changes are extremely small, on the order of magnitude of 10 fT [290]. Thus, MEG scanners need to be housed in magnetically shielded rooms (also known as Faraday cages). This type of shielding eliminates artefacts from the Earth's magnetic field and other external magnetic noise generators [162]. The changes in magnetic flux can be detected with a superconducting sensor coil, which triggers a superconducting current flow. This flux transformer is connected to an input coil which is in turn inductively coupled to a SQUID [162]. The SQUID (superconducting quantum interference device) is a fundamental building block of the modern MEG scanners. They allow converting the extremely small changes in magnetic field into voltage changes [162]. To improve the detection of artefacts, a standard flux transformer consisting of a gradiometer is built from two coils oriented one below another, horizontally to the scalp. In the case of the magnetic signals generated within the brain, the small magnetic field will be detected only by the coil closest to the scalp, as the magnetic field density decreases sharply with distance. In case of an artefact from, e.g. a magnetic field generated by a passing car, the distortion in the magnetic field will be detected simultaneously by the two

coils. Only asynchronous signals will be passed from the gradiometer to the amplifier and then to the analogue to digital transformer [162, 163].

Compared with EEG, MEG devices are much bigger, require liquid helium for cooling, and a special environment to be kept in; it follows that they are much more expensive [163]. MEG typically consists of a large tube, placed vertically above the head. However, recent years have seen the development of portable MEG devices based on optically pumped magnetometers [397]. These devices, however, despite bringing initial practical results, are still to be formally tested for scientific applicability [397]. Also, MEG scanners can detect only signals from tangentially oriented sources, because radial dipoles give rise to no external magnetic field [163]. An EEG device allows detecting cortical activity from both tangential and radial sources.

Unlike electrical signals, the magnetic fields originating from neural activity are hardly affected by the conductivity of the tissues within the scalp. Hence, the signal-to-noise ratio of MEG is much higher, which improves for instance the spatiotemporal localisation of the signals' sources [163]. Related to this, the time of preparation of a participant for the scanning tends to be much shorter in MEG. EEG devices require placement of a gel or a liquid on the participant's scalp to improve the electrodes' impedance [290]. This is often linked with a tedious and time-consuming removal of skin fat and hair, to make sure that the electrode connects to the scalp.

On the other hand, both EEG and MEG suffer from poor spatial resolution. For instance, fMRI devices allow detecting the activity of in-depth brain regions, like basal ganglia, or insula [37]. The contribution of electrical potentials, or magnetic fields is much weaker to M/EEG recordings. Furthermore, spatial distortion in EEG and MEG arises from differences in the electrical resistance of the head tissue [293]. MEG signals are less affected by this and have better spatial resolution than EEG (by a factor of around two) [162]. Both EEG and MEG allow, however, estimating an approximate location of cortical sources, but often only when additional steps are made (like structural MRI scanning).

Despite these limitations, both EEG and MEG have been extensively used in clinical and academic applications [163]. Typical features of magneto-encephalographic recordings analysed in the studies include: frequency band contributions, effective functional

connectivity analysis, independent components, energy of the signal in time, but also event related signals [293, 290, 163, 162]. For example, Event-Related Potentials (ERP) in the case of EEG, or Event-Related Fields (ERF) for MEG, are fast changes of amplitude triggered by an event [83]. A significant deviation (with respect to the baseline) from a signal averaged over multiple realisations of a task can be detected when a participant attends to a particular event in a sequence. These effects have become a major focus of clinical and cognitive psychology experiments, but have also been used in modern brain–computer interface implementations [232].

Recently, a number of studies have started using neural correlates detected by EEG, or MEG, to support theoretical modelling [130, 288, 289]. For example, in [289], a model was used to track the onset of the accumulation of evidence. It was shown that a type of ERP, the N200 peak, tracks non-decision times estimated by cognitive models, e.g. DDM [289]. These ERPs were found in the approximate area of the extrastriate cortex. The results of a similar modelling approach with single-trial ERP magnitude changes will be presented in Chapter 2 to inform a cognitive model of decision-making. As another example, [213] showed the benefits of using rapidly sampled EEG and MEG for a special type of generative model: the Dynamic Causal Model (DCM) [132]. The DCM in the M/EEG setting allows inverting a full spatiotemporal model of ERPs depending on the condition of the experiment. In this way, it is possible to invert the generative DCM, which provides conditional densities on the model parameters. These parameter estimates enable performing hypothesis testing under the data and model constraints [213, 132].

1.6 The decision-making process

In this thesis (particularly, in Chapters 2, 4 and 5), I focus on the modelling of the cognitive phenomenon of decision making. In this section, the most recent developments in this area of research are briefly summarised. For a detailed overview of decision making, one can refer to, e.g. [148].

The study of decision making is ubiquitous in many fields of research, such as psychology, economics, engineering (e.g. quality control), and political science [148, 362].

This shows how crucial is the understanding of the problem. Citing Mark Twain [411],

> *I must have a prodigious quantity of mind; it takes me as much as a week sometimes to make it up.*

we may identify the basic principle that a decision takes time that is needed to commit to a choice [148].

There are several types of decision. *Perceptual* decisions are made at a non–deliberative level and have an arbitrary criterion that defines the quality of a choice [104, 49, 470]. *Value-based* (or preferential) decisions are considered to be at a higher cognitive processing state. The quality of such decisions depends on the decision maker's subjective goals [104]. According to the [318], a decision process consists of five aspects: an option representation, valuation, action selection, outcome evaluation, and learning. In the projects presented in this thesis, I will focus only on the phases of valuation and action selection. For simplicity, it is often assumed that the choice is binary, meaning that there are only two options available [362, 276].

In Chapters 2 and 4, I focus on perceptual decision making. This field has been influenced by the experiment performed by Newsome, Britten, and Movshon [281]. They recorded spiking activity from neurons in extrastriate area MT/V5 of rhesus monkeys while those monkeys performed a moving dots direction discrimination task. First, the fidelity of the single-neuron response to the motion rivalled the fidelity of the monkey's behavioural reports [281]. Later, they observed that the saccade evoked by stimulation of the frontal eye field showed a systematic deviation in the direction of a preferred choice. This provided the initial evidence that the decision process, and not just its outcome, seems to occur in the brain circuitry that governs the movement response [351].

Apart from the neuronal evidence, the study of decision making has benefited from a branch of statistics known as *Signal Detection Theory* (SDT) [154, 49, 148], which is a general framework for interpreting data from experiments, in which accuracy is measured. For example, when two or more classes of stimuli are sampled repeatedly and an observer must select a response corresponding to the class actually presented, the accuracy in such tasks is limited by the sensitivity of the observer, which depends

on the degree of overlap between the distributions of a decision variable produced by the classes of stimuli [154]. In SDT, a raw representation of a piece of evidence gives rise to a decision variable (DV), upon which the brain applies a *decision rule* to make a choice. In the classic approach to SDT, the DV is a simple transformation of the sensory data that satisfies the weak constraint of being monotonically related to a likelihood, namely, the probability of observing this value, given a state of the world. Then, the decision rule is effectively a comparison with a criterion [154, 362]. In a motion detection experiment, for instance, the decision is based on the comparison of spike counts from two pools of neurons that are the most sensitive to the two directions of the motion [147].

However, SDT does not provide an explanation for the amount of time needed to complete a decision [362]. Hence, SDT is often paired with a *sequential sampling* [148], which finds how the evidence has been integrated over time to a threshold level, when the decision terminates with a choice. Nonetheless, the accumulation of decision via an integration is not the only way for a decision to emerge [148, 362].

The framework of perceptual decisions arising via the accumulation of evidence over time is the core of many models of reaction time [351, 50, 470, 289]. Under the SDT framework, perceptual decisions can be modelled in a number of ways that have been shown to converge to only two: drift-diffusion models (DDM) and linear ballistic accumulators (LBA) [414, 49]. Both approaches introduce the notion of a decision variable that accumulates evidence towards one of the options. In the DDM this variable represents the difference between evidences in favour of each option [325], and in the LBA model it is represented by two (for the binary case) independent accumulators that compete ('race') towards the threshold [58]. Typically, a decision model consists of the following parameters: the non-decision time (representing sensory information encoding and movement), drift rate or accumulation rate (indicating a force with which a decision variable is attracted towards an option), a threshold or boundary (representing the stopping point for the information accumulation), and bias (or a preference for one option) [49]. It has been shown that the drift-diffusion model behaves optimally when its processing remains in a linear range [49]. By 'optimally' we mean here the best theoretically possible performance in the sequential probability

ratio test (SPRT) sense [432].

Cognitive models are used to explain the general mechanism of decision making, but they do not answer the question of where to look in the brain for the neural correlates of a decision variable [297]. One of the first pieces of empirical evidence suggested that the lateral intraparietal area (LIP) might be engaged in this process [364, 362]. The LIP is part of Brodmann area 7, heavily connected with a brain structure involved in the control of eye movements [237]. Also, invasive recordings from the medial temporal (MT) area (involved in motion processing) from monkeys detected neuronal activity providing noisy evidence supporting two alternatives [53, 364]. These studies indicate that the choice is made when the firing rate of a certain neuronal pool reaches a decision threshold [364, 336, 160]. But are the decision variables the only parameters impacting the decision? Recent modelling work suggests that the decision threshold can be squeezed as an inward function of time [396]. This may be achieved by adding a time-dependent signal to the accumulated evidence, which we refer to as an 'urgency' signal [75]. With the more complex models of decision and inconclusive experimental work, currently, it is hypothesised that the LIP is one of many areas involved in decision making [207]. Other studies have mentioned the involvement of the motor and frontal areas of the brain [362, 160].

This short overview only touches the surface of contemporary research into decision making, and provides a brief background for the further overviews presented in Chapters 2, 4 and 5. Current research has arrived at the point where scientists are trying to bring together the three pillars of choice: accuracy, reaction time, and confidence, but also their neuronal correlates. This work is important, as decision making is one of the key components of human cognition [362].

1.7 Research questions and outline

In this thesis, I combine mathematical modelling, numerical simulations, and data analysis, to understand selected cognitive mechanisms of healthy and diseased human brains. I start from a high-level cognitive model of decision making. In Chapter 2, I combine EEG and cognitive models to understand the phenomenon of breaking a

deadlock during the selection of options with the same probability of reward. In Chapter 3, I use a data-driven model to find statistical differences of resting-state networks dynamics between juvenille myoclonic epilepsy patients and healthy controls. Next, in Chapter 4, I examine a more biologically-plausible model of perceptual decision making that explains the human performance when deciding on the basis of multiple information sources. Finally, in Chapter 5, I present a spiking neural network model with theoretically motivated connections fitted with NEF. Overall, this thesis covers a non-exhaustive but wide range of neuronal computational models. In Chapter 6, I summarise the key findings of the presented research and discuss its limitations.

Chapter 2

High-level model of cognition informed by EEG data: Reward certainty in the voluntary decision-making task

This chapter is based on the work published in the *Computational Brain & Behavior* [465] in collaboration with Wojciech Zajkowski (behavioural data analysis, modelling, interpretation of the results), Jacopo Barone (data collection, data pre-processing), Dr Lisa H. Evans (supervision of data collection and data pre-processing) and Dr Jiaxiang Zhang (experimental design, data analysis, methodology, general supervision). The author's contributions to this work include EEG data pre-processing and analysis, MVPA analysis, data visualisation, modelling and writing.

In this research, we aim to characterise the neurocognitive processes underlying voluntary decisions, i.e. when the outcome reward between two options is the same. We use hierarchical Bayesian parameter estimation for LBA model. We showed that the probability and preference effects were associated with changes in the speed of evidence accumulation, but not with visual encoding or motor execution latencies. Furthermore, we integrate cognitive modelling of behavioural responses with single-trial EEG features. The modelling showed that the rate of change between N100 and P300 event-related potentials modulated accumulation rates on a trial-by-trial basis.

2.1 Background

Cognitive flexibility enables decision strategies to adapt to environmental and motivational needs [352]. One characteristic of this ability is that harder decisions often take longer. Evidence from neurophysiology [147], neuroimaging [171] and modelling [327] suggest an evidence accumulation process for decision-making: information is accumulated over time, and a decision is made when the accumulated evidence reached a threshold [148]. This process can accommodate paradigms consisting of noisy stimuli (perceptual choices), as well as a rich variety of tasks with unambiguous stimuli (value-based [308] or memory-based choices [321]). For perceptual choices, evidence is derived from the sensory properties of the stimuli; for value or preference-based choices, it originates from internal value evaluation and comparison [226]; while for memory-dependent choices, from sampling memory traces [321, 363]. According to this framework, decision difficulty, and in turn response time (RT), is proportional to the relative difference in the evidence supporting each option, consistent with results from perceptual [99], value-based [310, 296] and memory based decisions [325].

Recent studies propose that the nervous system does not only estimate the evidence of options, but also multiplies it with a gradually growing 'urgency' signal [396]. Such approach represents a decision policy for maximising reward rate while taking into account the information conveyed by successive samples of the environment. It has been shown that higher reward rates are achieved when the accumulated evidence is novel, and the results are compared with a decreasing accuracy criterion [98]. However, the urgency gating can be a good explanation of the deadlock break only in a temporal dimension, but does not provide an explanation of the accuracy bias observed in this and other studies [69].

Making difficult choices requires more evidence, and hence longer deliberation can be an advantageous decision strategy. Scaling deliberation with difficulty is beneficial only to a certain point. What happens if decision difficulty reaches a tipping point with values of options being indistinguishable? In the hypothetical paradox of *Buridan's ass* [417], a donkey which cannot choose between two identical haystacks would, as a result of its indecision, starve to death. This view is consistent with the classical DDM [325],

which encodes the relative difference of evidence in favour of two options as a single accumulation process between two absorbing boundaries. Such a model would predict a deadlock or indecision between two equal alternatives, because there is zero difference in the mean evidence supporting each choice (e.g. two identical haystacks), and the decision process is dominated by noise accumulated over time, resulting in prolonged RT [391] (but see [329] for a recent model modification that addresses this theoretical limitation).

On the other hand, economic analysis suggests that choices between equal alternatives should be made as fast as possible. The benefit of 'rushing to decisions' comes from being able to relocate our cognitive resources elsewhere [342]. If evidence cannot bring us closer to a better choice, deliberative thinking becomes an expensive and unnecessary luxury. This effect can be modelled using stochastic decision models with multiple accumulators, each encoding the accumulated evidence in favour of one choice, such as the Linear Ballistic Accumulator (LBA) model [58] and the Leaky Competing Accumulator (LCA) model [414, 50]. For those models, multiple accumulators compete against each other on the basis of multiple sources of evidence inputs, which by default eliminates the scenario of indecision between equal alternatives.

In reality, individuals can make timely choices between equally valued options [424] and their decision speed is facilitated by the level of reward magnitude [307]. For example, in preference-based decisions, it took under 2 seconds for one to choose between two snack food stimuli that had similar valuations [424]. In both humans and non-human primates, higher reward magnitude facilitates RT in perceptual and value-based decisions between equal choices [307]. Intuitively, Buridan's donkey would be motivated to make faster decisions if the haystacks are fresh, compared to when they are stale. This magnitude effect is in line with ecological incentives: high rewards may imply a resource-rich environment, for which one needs to exploit as early as possible; low rewards may imply a resource-poor environment in which it is worth waiting for a better option [307]. Furthermore, if choices are based purely on expected rewards, one may choose any of the equal-valued options with the same frequency, leading to random behavior. Nevertheless, previous studies [471, 305] showed that in a sequence of voluntary action decisions, humans deviated from a random pattern of choice and

exhibited low choice entropy across trials. A similar conclusion has been reached in consumer decisions, where brand loyalties are driven by seemingly irrational preferences [440]. These findings suggest a possible preference bias between equal options, which renders some options more likely to be chosen than others.

We focus on three issues that have been unresolved in previous research on choices between equal alternatives. First, we aim to explore the effect of reward probability on RT. We expect that, similar to magnitude [391, 307], higher reward probability accelerates RTs. This prediction is not trivial, since probability and magnitude can have different effects on behaviour. For example, [463] showed that magnitude discounting follows a power law, while probability is discounted hyperbolically. Unlike magnitude, probability has a upper bound at 100%, which acts in a qualitatively distinct way on behaviour [409]. We expect this increase in speed to be non-linear, with choices between two certain (100% probability) options being disproportionately faster compared to choices between two uncertain ones.

Second, in the evidence accumulation framework, both the rate of the accumulation and the non-decision time can influence a model's prediction of reaction time, the former encoding the strength of evidence and the latter reflecting the latencies of visual encoding and motor execution. During perceptual learning, the accumulation rate increases along with behavioural improvements [198], while the non-decision time remains unchanged in the late stage of training [470]. Furthermore, the accumulation rate is associated with the individual differences in working memory [354] and attention [288], while the non-decision time is faster in individuals with higher diffusion MRI derived neurite density in the corticospinal tract, the primary motor output pathway [206]. Recent research showed that both parameters can be influenced by reward magnitude [430], and the current study will examine further whether reward probability and preference influence the two model parameters.

Third, we aim to describe the macroscopic pattern of brain activities associated with differences in behaviour: it's temporal evolution and relation to model-derived parameters. Functional imaging studies have localised the mesocorticolimbic dopaminergic network to be involved in both reward certainty and preference processing [401, 1], but little is known about how these relate to global activations across the scalp. Pin-

pointing when EEG activity diverges between conditions and assessing whether these differences are transient or sustained can further inform our computational model, giving deeper insight into the cognitive underpinnings of the decision process.

Here, we address these questions by combining advanced computational modelling and EEG in a probabilistic reward task. Participants memorised six unambiguous cues associated with three levels of reward probability, a certain reward level (i.e. 100%) and two levels of uncertain reward probabilities (80% and 20%). Participants made two-alternative forced choices between cues with equal reward probability (Figure 2.1). The inclusion of the 100% reward probability condition allowed us to investigate whether cues with definitive rewards are processed in a different manner than the uncertain cues [118]. Additional task conditions involved binary decisions between cues with different reward probability (unequal trials) and unitary responses to single cues (single-option trials). This design enabled us to focus on the neurocognitive processes underlying choices between equal options, while participants maintained a clear understanding of cue values for rational decisions between unequal options.

We first examine how reward probability influences behaviour and whether a preference bias between equal options is present. We then fit an accumulator model of decision-making [58] to the behavioural performance across reward probability levels. Posterior group parameters from hierarchical Bayesian model fitting procedure were used to infer whether the behavioural effects were driven by evidence accumulation or non-decisional components of the process. EEG data were analysed with time-resolved multivariate pattern classification for decoding spatio-temporal representations of reward probability and preference. To establish a link between the decision process and its EEG signatures, we combined behavioural and EEG data into a joint hierarchical Bayesian model and tested the hypothesis that electrophysiological activity reflects trial-by-trial changes in the speed of evidence accumulation for decisions [412].

We demonstrate that reward probability and spontaneous preference independently shape RTs and choices when deciding between equal alternatives. These behavioural effects affect the decision process and evoke a distinct electrophysiological pattern. Together, our findings contribute to the understanding of how decision deadlocks between two equally probable rewards can be overcome.

2.2 Materials and Methods

2.2.1 Participants

Twenty-three healthy participants were recruited from Cardiff University School of Psychology participant panel (20 females; age range 19-32, mean age 22.7 years; 22 right-handed). All participants had normal or corrected-to-normal vision, and none reported a history of neurological or psychiatric illness. Written consent was obtained from all participants. The study was approved by the Cardiff University School of Psychology Research Ethics Committee.

2.2.2 Apparatus

The experiment was conducted in a dedicated EEG testing room. A computer was used to control visual stimulus delivery and record behavioural responses. Visual stimuli were presented on a 24-inch LED monitor (ASUS VG248) with a resolution of 1920 by 1080 pixels and a refresh rate of 60 Hz, located approximately 100 cm in front of participants. Participants' responses were collected from a response box (NATA technologies). The experiment was written in Matlab (Mathworks; RRID: SCR_001622) and used the Psychophysics Toolbox Version 3 extensions [222].

2.2.3 Experimental design

All participants performed a decision-making task with probabilistic rewards during EEG recording (Figure 2.1). Before the task, participants memorised 6 unambiguous cues represented by different symbols and their associated probabilities of receiving a reward (Figure 2.1B; see Section 2.2.4). All the cues had the same colour (RGB = 246, 242, 92) on a black background (100% contrast). Each cue was mapped onto one of the three reward probability levels: *high* (a reward probability of 100%, i.e. always rewarded), *medium* (a reward probability of 80%) and *low* (a reward probability of 20%), and hence there were two different cues associated with each reward probability.

Participants were instructed to maximise the total accumulated reward in the decision-making task. The task contained three types of trials: equal, unequal and

single-option. On an equal trial, two different cues with the same reward probability appeared on the left and right sides of a central fixation point (e.g. 100% vs. 100%, 80% vs. 80% or 20% vs. 20%). On an unequal trial, two cues with different reward probability levels appeared on both sides of the central fixation point (e.g. 100% vs. 20%, 100% vs. 80% or 80% vs. 20%). On a single-option trial, one of the six cues appeared on either the left or right side of the fixation point. In equal and unequal trials, participants chose the left or right cue via button presses with the right-hand index and middle fingers. We limited a choice to the right hand in order to avoid the impact of handedness bias. In single-option trials, participants responded to which side the single cue was presented (i.e. left or right). In all trials, the reward was operationalised as 10 virtual 'game points' that did not have any tangible value. The probability of receiving the reward in a trial was either 100%, 80% or 20%, which was determined by the chosen cue. It is worth noting that, in equal trials, participants' decisions did not actually affect the probability of receiving the reward because both options had equal reward probability. In single-option trials, if participants chose the wrong side with no cue presented (0.1% across all single-option trials), no reward was given. Feedback for rewarded (a '10 points' text message on the screen) or not rewarded (blank screen) choices was given after each trial. The total game points awarded were presented at the bottom of the screen throughout the experiment.

2.2.4 Procedure

Each experimental session comprised 640 trials, which were divided into 4 blocks of 160 trials. Participants took short breaks between blocks and after every 40 trials within a block. The mapping between the six reward cues and three levels of reward probability was randomised across participants. During breaks, the cues-reward mappings were explicitly presented on the screen (Figure 2.1B), and the participants could take as much time as they needed to memorise them. After the first two blocks, all the cues were re-mapped to different reward probabilities. For example, for the pair of two cues that were associated with 100% reward probability in the first and second blocks, one of the two cues would be associated with 80% reward probability in the third

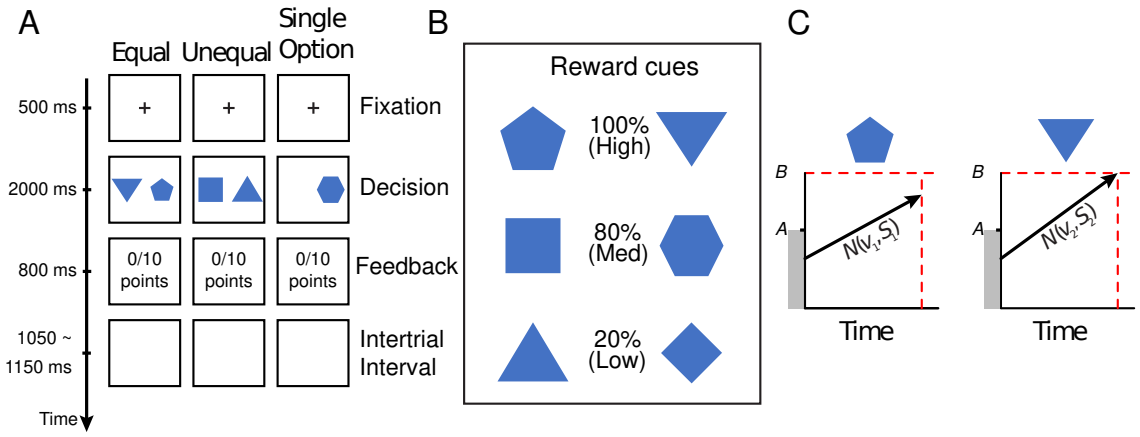

Figure 2.1: *(A) Experimental paradigm of the probabilistic reward task. Participants were instructed to decide between two reward cues (equal and unequal trials) or respond to a single cue (single-option trials). (B) A total of six reward cues were randomly assigned to three levels of reward probability (100%, 80% or 20%). (C) Exemplar time course of the Linear Ballistic Accumulator (LBA) model for equal choices. On each trial, the LBA assumes that evidence for two options are accumulated linearly and independently over time in two accumulators. The accumulation rate is sampled from a normal distribution with mean v and standard deviation S. The starting point of the accumulation process is sampled from a uniform distribution between 0 and A. The accumulation process terminates once the accumulated evidence first reaches a threshold B, and a corresponding decision is made by the winning accumulator.*

and fourth blocks, and the other associated with 20% reward probability. Participants were encouraged to memorise the altered cue-probability associations prior to the third block. This remapping procedure reduced the potential bias associated with specific cues. No explicit memory tests were performed.

Each block contained 64 equal trials (32 for 100% vs. 100%, 16 for 80% vs. 80% and 16 for 20% vs. 20%); 64 unequal trials (32 for 80% vs. 20%, 16 for 100% vs. 80% and 16 for 100% vs. 20%) and 32 single-option trials (16 for 100%, 8 for 80% and 8 for 20%) at a randomised order. This design ensured the same number of trials with and without cues with the highest reward probability (100%). Note, however, that individual cues did not differ much in terms of frequency of occurrence: each 100% cue appeared 56 times, compared to 48 for each non-certain cue. This makes it unlikely that observed differences can be explained by occurrence frequency alone. Because two cues were bound to every probability level, different cue positions and combinations can result in the same reward probability pair (e.g. there are 4 possible combinations for 80% vs. 20% unequal trials). These combinations were counterbalanced across trials.

Each trial began with the presentation of a fixation point at the centre of the screen for 500 ms. After the fixation period, in the equal and unequal trials, two reward cues appeared on the left and right sides of the screen with a horizontal distance of 4.34° from the fixation point. Both cues were vertically centred. In single-option trials, only one reward cue appeared on one side of the screen, and the side of cue appearance was randomised and counterbalanced across trials. Cues were presented for a maximum of 2000 ms, during which participants were instructed to make a left or right button press. The cues disappeared as soon as a response was made, or the maximum duration was reached. The reaction time (RT) on each trial was measured from the cue onset to button press. Reward feedback was given 200 ms. after the reward cue offset and lasted 800 ms, followed by a random intertrial interval uniformly distributed between 1050 and 1150 ms. As in the previous study [470], if the participant failed to respond within 2 sec or responded within 0.1 sec, no reward was given and a warning message ‘too slow’ or ‘too fast’ was presented for 1.5 sec.

2.3 Data analysis

2.3.1 Behavioural analysis

We excluded trials with RT faster than 200 ms (fast guesses). For each participant, trials with RTs longer than 2.5 standard deviations from the mean RT were also excluded from subsequent analysis. The discarded trials accounted for 1.5% of all trials.

We first analysed the proportion of choices in equal trials to establish the existence of a preference bias. In the equal condition, by definition, there was no 'correct' or 'incorrect' response, since the cues had the same reward probability. For each pair of cues with the same reward probability, we defined the preferred cue as the one chosen more frequently than the other (non-preferred) in equal trials. The categorisation of preferred and non-preferred cues was estimated separately between the first two and the last two blocks, because of the cue-probability remapping after the first two blocks. At each level of reward probability, a preference bias was then quantified as the proportion of trials where the preferred cue was chosen. The preference bias had a lower bound of 50%, at which both cues were chosen with equal frequency.

In the unequal condition, we defined decision accuracy as the proportion of choosing the cue with higher reward probability, separately for each combination of reward probabilities (100% vs. 80%, 100% vs. 20% and 80% vs. 20%). Two-tailed one-sample t-tests compared the decision accuracy in the unequal condition against a chance level of 50%, which would indicate irrational decisions (i.e. both high and low reward cues were chosen in 50% of trials).

To determine how reward probability, preferences and other experimental factors influence RT, we analysed single-trial RT data with linear mixed-effects models (LMM) using the `lme4` package [26] in *R* (RRID: SCR_001905). The LMM is a hierarchical regression method that distinguishes between fixed and random effects [156]. LMMs take into account all single-trial data without averaging across trials and offer better control of type 1 and type 2 errors than ANOVA [16]. Therefore, statistical inferences from LMMs are robust to experimental designs with unbalanced trials across conditions [19], which is an important feature suitable for the current study.

We designed two LMMs with different dependent variables and factors (Table 2.1). Model 1 analysed the RTs from equal and single-option trials, including choice type (equal or single-option), reward probability (high, medium or low), cue remapping (before and after), preference (whether the chosen cue was preferred) and right-side bias (whether the chosen cue was on the right side of the screen) as factors. Right-side bias was included to control for spatial bias relating to preference for stimuli presented on the right or left side of the screen. For the unequal condition, because each trial had two cues with different levels of reward probability that cannot be directly compared with equal or single-option trials, the RTs were analysed separately in Model 2. Here, we used similar predictors with exception of probability, which was captured by two additional factors: the sum and the absolute difference of the two reward probabilities, as they both have been shown to affect choice behaviour [393, 22, 391].

In all the LMMs, fixed effects structures included hypothesis-driven, design-relevant factors and their interactions, and individual participants were included as the source of random variance (random effect). We used a standard data-driven approach to identify the random effects structure justified by the experimental design, which resulted in good generalisation performance [24]. This approach starts with the maximal random effects structure (i.e. including all random slopes, intercepts and interactions) and systematically simplified it until the LMM reaches convergence. Table 2.1 lists the simplified random effects structures. The correlation structures of each fitted LMM was assessed to avoid overfitting [263].

2.3.2 A Cognitive Model of Voluntary Decision-Making

We further analyzed the behavioral data using the Linear Ballistic Accumulator (LBA) model [58]. LBA model is a simplified implementation of a large family of sequential sampling models of decision-making [327, 49, 148, 468] which assumes an independent accumulation process for each choice option. Our model-based analysis has three stages. First, we fit a family of LBA models with various model complexity to the behavioural data of individual participants in equal trials. By identifying the best-fitting model, we infer how reward probability and preference modulated subcomponents of

Table 2.1: *The linear mixed-effects models of RT. Model 1 analysed single-trial RT in equal and single-option trials. Model 2 analysed single-trial RT in unequal trials. In both models, preference was a predictor indicating whether the preferred cue was selected in each trial. Cue-remapping was a predictor indicating whether each trial was before or after cure-probability remapping in the second half of each session. Right-bias indicated whether the cue on the right size of the screen was chosen in each trial, modelling a possible response bias.*

	Model 1	Model 2
Dependent Variables	RT	RT
Main Effects	reward probability preference cue-remapping choice (equal of single-option) right-bias	sum of reward probability difference of reward probability preference cue-remapping right-bias
Interaction Terms	probability * choice probability * preference probability * cue-remapping choice * preference choice * cue-remapping preference * cue-remapping probability * choice * preference probability * choice * cue-remapping probability * cue-remapping * preference	sum of reward probability * preference difference of reward probability preference sum of reward probability * cue-remapping difference of reward probability * cue-remapping preference * cue-remapping
Random Effects (correlated slopes) and intercepts)	reward probability preference cue-remapping choice right-bias	sum of reward probability difference of reward probability preference cue-remapping right-bias

the evidence accumulation process during decision-making. Next, we simulate the best fitted LBA model and examine whether model simulations are consistent with the experimental data in single-option and unequal conditions. This is a stringent test of model generalisability because the experimental data in single-option and unequal trials are unseen by the model fitting procedure. Finally, we link the cognitive processes identified by the LBA model to brain activities by incorporating a trial-by-trial measure of EEG activity regressors into the best-fitted model [67, 288, 289].

The LBA model assumes that the decision of when and which to choose is governed by a 'horse race' competition between two accumulators $i \in \{1, 2\}$ that accumulate evidence over time supporting the two choice options. One accumulator is in favour of the preferred cue and the other of the non-preferred cue. The activations of the accumulators represent the accumulated evidence. At the beginning of each trial, the initial activation of the two accumulators are independently drawn from a uniform distribution between 0 and A. The activation of each accumulator then increases linearly over time, and the speed of accumulation (i.e. accumulation rate) varies as a Gaussian random variable with mean v_i and standard deviation S_i across trials. The accumulation process terminates when the activation of any accumulator reaches a response threshold B ($B > A$) and the choice corresponding to the winning accumulator is selected. The model prediction of RT (measured in seconds) is the sum of the duration of the accumulation process and a constant non-decision time T_{er}, with the latter accounts for the latency associated with other processes including stimulus encoding and action execution [58, 289, 206].

2.3.3 Model Parameter Estimation and Model Selection

LBA model has five key parameters: mean v and standard deviation S of the accumulation rate across trials, decision threshold B, starting point variability A and non-decision time T_{er}. To accommodate the empirical data, one or more model parameters need to vary between conditions. We evaluated a total of 21 variants of the LBA model with different parameter constraints (Figure 2.3A). First, the accumulation process may differ between the preferred and non-preferred options, leading v or S

to vary between accumulators (preferred, non-preferred). Second, reward probability could modulate the accumulation process or visuomotor latencies unrelated to decisions, leading to v, S or T_{er} to vary between three levels of reward probability. Third, the decision threshold B and starting point A were fixed between conditions, because the trial order was randomised, and we do not expect the participants to systematically vary their decision threshold before knowing the cues to be presented [327]. Fourth, decision threshold B and starting point A were fixed across preference levels, since participants could not predict which cue would appear on which side of the screen. During model-fitting, the decision threshold was fixed at 3 as the scaling parameter [58], and all the other parameters allowed to vary between participants. Theoretically, the scaling parameter can be set to an arbitrary value, which does not influence the parameter inference, as long as the priors of other parameters remain realistic, but with some constraints parameter estimation is easier to converge. Finally, because the participants showed behavioural differences between reward probability levels and between preferred/non-preferred choices, we only estimated realistic models: those with at least one parameter varied between reward probability levels (v, S or T_{er}) and at least one parameter varied between accumulators (v or S).

We use a hierarchical Bayesian model estimation procedure to fit each LBA model variant to individual participant's choices (the proportion of preferred and non-preferred choices) and RT distributions in equal trials. The hierarchical model assumes that model parameters at the individual-participant level are random samples drawn from group-level parameter distributions. Given the observed data, Bayesian model estimation uses MCMC methods to simultaneously estimate posterior parameter distributions at both the group level and the individual-participant level. The hierarchical Bayesian approach has been shown to be more robust in recovering model parameters than conventional maximum likelihood estimation [195, 472].

For group-level parameters (v, S, A and T_{er}), similar to previous studies [10], we used weakly informed priors for their means $E(.)$ and standard deviations $std(.)$:

$$
\begin{aligned}
E(v) &\sim N(2.5, 1), & std(v) &\sim \gamma(1, 1), \\
E(S) &\sim N(1, 0.75), & std(S) &\sim \gamma(1, 1), \\
E(A) &\sim N(2.5, 1), & std(A) &\sim \gamma(1, 1), \\
E(T_{er}) &\sim N(0.5, 0.2), & std(A) &\sim \gamma(1, 1),
\end{aligned}
\tag{2.1}
$$

where N represents a positive normal distribution (truncated at 0) with parameters mean and standard deviation, and γ represents a gamma distribution with parameters mean and standard deviation.

We used the `hBayesDM` package [4] in R for the hierarchical implementation of the LBA model. `hBayesDM` uses `Stan` for Bayesian inference and offers a range of hierarchical Bayesian implementations of various computational models on an array of decision-making tasks. For each of the 21 model variants, we generated four independent chains of 7,500 samples from the joint posterior distribution of the model parameters using Hamiltonian Monte Carlo (HMC) sampling in Stan [66]. HMC is an efficient method suitable for exploring high-dimensional joint probability distributions [38]. The initial 2,500 samples were discarded as burn-in. To assess the convergence of the Markov chains, we calculated Gelman-Rubin convergence diagnostic $\hat{R}$ of each model [138] and used $\hat{R} < 1.1$ as a stringent criterion of convergence [10]. We compared the fitted LBA model variants using Bayesian leave-one-out information criterion (LOOIC). LOOIC evaluates the model fit while considering model complexity, with lower values of LOOIC indicating better out-of-sample model prediction performance [420].

2.3.4 EEG data acquisition and processing

EEG data were collected using a 32-channel Biosemi ActiveTwo device (BioSemi, Amsterdam). Due to technical issues, EEG data collection was not successful in two participants, and therefore all EEG data analyses were performed on the remaining 21 participants. EEG electrodes were positioned at standard scalp locations from the International 10-20 system. Vertical and horizontal eye movements were recorded using bipolar electrooculogram (EOG) electrodes above and below the left eye as well as from the outer canthi. Additional electrodes were placed on the mastoid processes.

EEG recordings (range DC-419 Hz; sampling rate 2048 Hz) were referenced to linked electrodes located midway between POz and PO3/PO4 respectively and re-referenced off-line to linked mastoids. Additional electrodes were placed on the mastoid processes. EEG (range DC-419 Hz; sampling rate 2048 Hz) was collected with respect to an active electrode (CMS; common mode sense) and a passive electrode (DRL; driven right leg), which were located midway between POz and PO3/PO4 respectively, to form a ground-like feedback loop.

EEG data were pre-processed using EEGLab toolbox 13.4.4b [91] in Matlab 2016b. The raw EEG data were high-pass filtered at 0.1 Hz, low-pass filtered at 100 Hz using Butterworth filters and downsampled to 250 Hz. An additional 50 Hz notch filter was used to remove mains interference. We applied Independent Component Analysis (ICA) to decompose continuous EEG data into 32 spatial components, using `runica` function from the EEGLab toolbox. Independent components reflecting eye movement artefacts were identified by the linear correlation coefficients between the time courses of independent components and vertical and horizontal EOG recordings. Additional noise components were identified by visual inspection of the components' activities and scalp topographies. Artefactual components were discarded, and the remaining components were projected back to the data space.

After artefact rejection using ICA, the EEG data were low-pass filtered at 40 Hz and epoched from -400 ms to 1000 ms, time-locked to the onset of the stimulus (i.e. reward cues) in each trial. Every epoch was baseline corrected by subtracting the mean signal from -100 ms to 0 ms relative to the onset of reward cues.

2.3.5 Multivariate pattern analysis

We use time-resolved Multi-Voxel Pattern Analysis (MVPA) on pre-processed, stimulus-locked EEG data to assess reward-specific and preference-specific information throughout the time course of a trial. In contrast to univariate ERP analysis, MVPA combines information represented across multiple electrodes, which has been shown to be sensitive in decoding information representation from multi-channel human electrophysiological data [73, 97].

We conduct three MVPA analysis to identify the latency and spatial distribution of the EEG multivariate information. The first to decode reward probability levels in equal choices (e.g. equal trials with two 100% reward cues versus equal trials with two 80% cues). The second to decode preferred versus non-preferred choices in equal trials. The third to decode between equal and single-option choices with the same reward probability (e.g. equal trials with two 100% cues versus single-option trials with a 100% cue).

Each analysis is formed as one or multiple binary classification problems, and the data features for classification included EEG recordings from all 32 electrodes. In each analysis, at each sampled time point (-400 ms to 1000 ms) and for each participant, we train linear support vector machines (SVM) [135] using the 32-channel EEG data and calculate the mean classification accuracy following a stratified ten-fold cross-validation procedure. In all MVPA, we include the EEG data from 400 ms before cue onset as a sanity check, because one would not expect significant classification before the onset of reward cues.

In each cross-validation, 90% of the data issued as a training set, and the remaining 10% as a test set. In some analysis (e.g. equal trials with 100% cues versus equal trials with 80% cues), the number of samples belonging to the two classes is unbalanced in the training set. We use a data-driven over-sampling approach to generate synthetic instances for the minor class until the two classes had balanced samples [467]. The synthetic instances are generated from Gaussian distributions with the same mean and variance as in the original minority class data. Training set data were standardised with z-score normalisation to have a standard normal distribution for each feature. The normalisation parameters estimated from the training set was then applied separately to the test set to avoid overfitting. To reduce data dimensionality, we perform principal component analysis to the training set data and selected the number of components that explained over 99% of the variance in the training set. The test set data are projected to the same space with reduced dimensions by applying the eigenvectors of the chosen principal components. We then train SVM to distinguish between the two classes (i.e. conditions) and evaluate the classification accuracy using the test set data. The procedure is repeated ten times with different training and test sets, and

the classification accuracies are averaged from the ten-fold cross-validation. We use the SVM implementation in MATLAB Machine Learning and Statistics Toolbox. The trade-off between errors of the SVM on training data and margin maximisation is set to 1.

To estimate the significance of the classification performance, we use two-tailed one-sample t-test to compare classification accuracies across participants against the 50% chance level. To account for the number of statistical tests at multiple time points, we use cluster-based permutation [258] to control the family-wise error rate at the cluster level from 2000 permutations.

2.3.6 Estimation of single-trial ERP components

We estimate two ERP components from single-trial EEG data in equal trials: N100 and P300, which are subsequently used to inform cognitive modelling. The visual N100 is related to visual processing [256] and the P300 is related to evidence accumulation during decision making [209, 412].

To improve the signal-to-noise ratio of single-trial ERP estimates, we use a procedure similar to previous studies [208, 303, 289]. For each participant, we first performed singular value decomposition (SVD) to the grand averaged ERP data across all trials from the same experimental condition. SVD decomposes the trial-averaged ERP data $A_{k\times p}$ (where k is a number of channels and p is a number of time points) into independent principal components. Each component consists of a time series of that component and a weighing function of all channels, defining the spatial distribution (or spatial filter) of that component. Because the ERP waveform is the most dominant feature of the trial-averaged ERP data, the time course of the first principal component (i.e. the one that explains the most variance) represents a cleaned trial-average ERP waveform [289], and its weight vector provides an optimal spatial filter to detect the ERP waveforms across EEG channels. We then applied the spatial filter from the first principal component as a channel weighting function to single-trial EEG data to improve the signal-to-noise ratio.

The single-trial EEG data filtered with the SVD-based weighting function is then

used to identify the peak-latency and peak-amplitude of the N100 and P300 components. For N100, we search for the peak negative amplitude in a window centred at the group-level N100 latency (112 ms) and started at 60 ms. The lower bound of the search window was determined by the evidence that the visual onset latency is 60 ms in V1 [355]. For P300, we search for a peak positive amplitude in a window centred at the group-level P300 latency (324 ms). For both N100 and P300, the search window has a length of 104 ms, similar to a previous study [289].

Recent studies showed that the variability of the P300 component closely relates to the rate of evidence accumulation during decision making [412]. We therefore extend the best fitting LBA model with EEG-informed, single-trial regressors, which estimates the effect of trial-by-trial variability in EEG activity on the mean accumulation rate [169, 288].

The main regressor of interest is the slope of change between the N100 and P300 components, which is defined as the ratio of the P300-N100 peak-amplitude difference and the P300-N100 peak-latency difference in each equal trial. We also test four additional regressors from individual ERP components: P300 amplitude, P300 latency, N100 amplitude and N100 latency. All the EEG regressors are obtained from the estimations of single-trial ERP components in equal choice trials. To obtain a meaningful intercept, the regressors are mean-centered and rescaled to have a unit standard deviation.

Each EEG regressor is tested in a linear regression model, using the same Bayesian hierarchical model estimation procedure as in the behavioural modelling analyses. For each regression model, we assume that the mean accumulation rates of both accumulators $v_1(t)$ and $v_2(t)$ (i.e. the one in favor of the preferred option and the other one in favor of the non-preferred option) are influenced by the EEG regressor of interest on a trial-by-trial basis:

$$v_1(t) = \widetilde{v_1} + \beta \times EEG(t), \quad v_2(t) = \widetilde{v_2} + \beta \times EEG(t), \tag{2.2}$$

where $t = 1, 2, 3, \ldots$ represents the equal choice trials, and $\widetilde{v_1}$ and $\widetilde{v_2}$ are the intercepts. The regression coefficient β represents the effect of EEG regressor on the mean

accumulation rates.

The rationale of estimating an EEG regressor to the mean drift rate is twofold. First, this approach allows quantifying the trial-by-trial change over the intercept (i.e. the mean drift rate), independent of its trial-by-trial variability (parameter S). Second, one would not expect the sensor level EEG signal has sufficient spatial resolution to distinguish between the two accumulators encoding two options. Therefore, we estimated a single EEG regressor across both accumulators.

2.3.7 Open data and scripts

We have made the data (https://figshare.com/articles/dataset/9989552), and all analyses scripts (for modelling and EEG analysis) openly available (https://github.com/ccbrain/voluntary-decision-eeg).

2.4 Results

2.4.1 Behavioural results

For each pair of cues with the same reward probability, we defined the preferred cue as the one chosen more frequently than the other (non-preferred) in the equal choice trials (see behavioural Analysis in Methods). We found a strong preference bias (>50%) for choosing one reward cue over the other at each level of reward probability (Figure 2.2C; high: 95%CI = $[0.682, 0.765]$; medium: 95%CI = $[0.679, 0.759]$; low: 95%CI = $[0.669, 0.745]$). A repeated-measures ANOVA showed no significant difference in preference between reward probability levels ($F(2, 44) = 0.2, p = 0.81$). Therefore, although the two options were associated with the same level of reward probability, participants did not make their choices randomly. We further used a LMM to evaluate the preference bias as a function of cue remapping (i.e. before vs. after the cue remapping halfway through each session) and trial order in each testing block. The preference bias was smaller after cue remapping ($\beta = -0.181$, 95%CI = $[-0.01, -0.348]$, $p < 0.03$), but was not influenced by trial order ($\beta = 0.037$, 95%CI = $[-0.170, 0.243]$, $p = 0.73$). These results imply that, for a given set of cue-probability associations, the extent

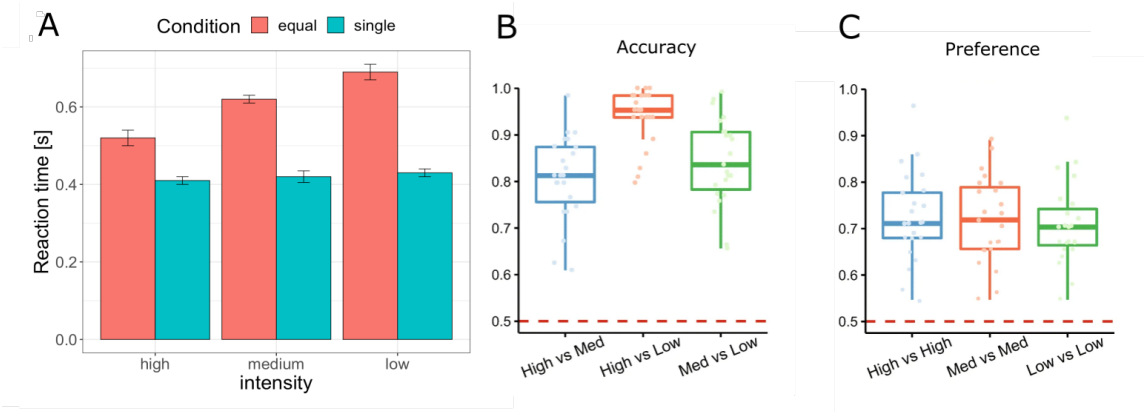

Figure 2.2: *(A) Reaction times between the conditions of the experiment. (B) Decision accuracy across reward probability levels in unequal trials. (C) Preference bias towards selected symbol.*

of preference bias did not significantly vary over time. Because the cue-probability mapping was randomised across participants and re-mapped within each session, the observed preference bias is unlikely to be explained by a group-level preference towards any specific cue, but rather a spontaneous preference at the individual level. Additionally, to check if preference from first half of the experiment affected preference after remapping, we calculated the proportion of any cue being preferred in both sessions on a subject level. We found that preference was consistent only in 51.5% of cases, rendering no support for preference transfer after remapping (one-sided binomial-test; $p = 0.5$, 95%CI $= [0.361, 1]$).

Response Times

We used a LMM to quantify the influence of experimental factors on RTs in equal and single-option choices. The fixed effects included reward probability, choice type (equal vs. single-option), preference (choosing the preferred vs. the non-preferred option), cue remapping and their meaningful interactions. Participants were faster when choosing the preferred than the non-preferred option ($\beta = -0.063$, 95%CI $= [-0.027, -0.991]$, $p < 0.05$) and RTs decreased as the reward probability increased ($\beta = -0.101$, 95%CI $= [-0.067, -0.135]$, $p < 0.001$). The RT in equal choice trials were longer than that in single-option trials ($\beta = -0.292$, 95%CI $= [-0.201, -0.384]$,

$p < 0.001$). The effect of reward probability on RT was stronger in equal compared to single-option choices, supported by a significant interaction between the two main effects ($\beta = 0.045$, 95%CI = $[0.025, 0.066]$, $p < 0.001$).

Participants had slower responses after memorising a new set of cue-probability associations, indicated by a significant main effect in RT before and after cue remapping ($\beta = 0.149$, 95%CI = $[0.096, 0.201]$, $p < 0.001$). The significant interaction between cue remapping and reward probability suggested that the increase in RT was more pronounced in trials with lower reward probability ($\beta = -0.039$, 95%CI = $[-0.051, -0.026]$, $p < 0.001$). The interaction between cue remapping and choice type ($\beta = -0.247$, 95%CI = $[-0.192, -0.302]$, $p < 0.001$) indicated that this pattern was mainly associated with equal trials. Because evaluating reward probability of a cue was likely associated with additional cognitive load after cue remapping, the observed RT difference before and after cue remapping implies that participants evaluated both cues throughout the experimental session.

2.4.2 Cognitive modelling of behavioural data

To identify the cognitive processes that led to the observed behavioural differences, we compared 21 variants of the LBA model. The model variants differed systematically in their constraints on whether the rate of evidence accumulation and non-decision time could change between reward probability levels or preferred/non-preferred options. For each model variant, we used hierarchical Bayesian modelling with Markov chain Monte Carlo (MCMC) parameter estimation routine to estimate the posterior distributions of the model parameters, given the observed choice and RT distribution from individual participants (see Section 2.3.3). To identify the model with the best fit, we calculated the Bayesian LOOIC score for each model [420].

MCMC chains representing posterior parameter estimates in all the 21 model variants reached high levels of convergence (Gelman-Rubin convergence diagnostic $\hat{R} \leq 1.02$ for all parameters in all models). The LOOIC scores suggested that the models with the mean accumulation rate varying between reward probability levels and between preference levels fitted the data better than others model variants. The best-fitting

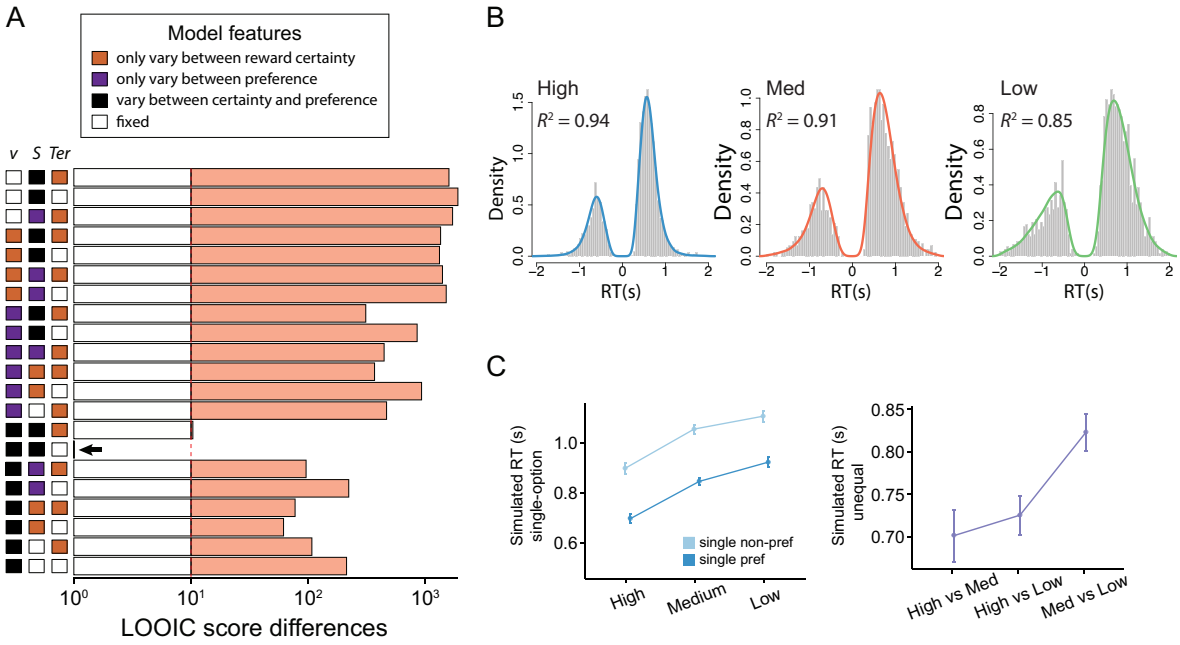

Figure 2.3: *Model comparisons, model fits and model simulations. (A) LOOIC scores of 21 LBA model variants. The LOOIC score differences between all models and the best model are plotted against corresponding model structures, which were illustrated on the left of the figure. The model structure specified how the mean accumulation rate v, the standard deviation S of the accumulation rate and the non-decision time T_{er} could vary between conditions. A black filled square indicated that the corresponding parameter could vary between reward probability levels and preferred/non-preferred options. An orange or purple filled square indicated that the corresponding parameter could only vary between reward probability levels or preferred/non-preferred options, respectively. Unfilled (white) squares indicated that the parameter remained fixed between conditions. Bar colour indicates whether the difference in LOOIC scores is considered substantial (over 10): white part of the bar corresponds to score up to 10, orange to the amount exceeding 10. The best model was shown with a LOOIC score difference of zero (indicated by the red arrow). (B) Simulations of RTs in equal choices, generated from the posterior distribution of the best fitted model for high (left), medium (middle) and low (right) reward probability levels. Grey histograms represent experimental data and solid coloured lines represent model simulations from 100 iterations for high (blue), medium (red) and low (green) conditions. Negative values represent RTs for non-preferred choices. (C) Testing the model on the unseen conditions. Simulated RTs in single-option condition (left) and unequal choices (right) from 100 iterations. Error bars represent standard errors across participants.*

model (i.e. the one with the lowest LOOIC score, Figure 2.3A) had fixed group-level non-decision time with the standard deviation of the accumulation rate varying between reward probability levels and preferred/non-preferred options. To evaluate the model fit to the empirical data in equal trials, we calculate the posterior prediction of the best fitting model by averaging 100 iterations of model simulation using posterior parameter estimates. Averaging across multiple iterations reduces potential biases when sampling from posterior parameter estimates. Each of the 100 iterations generates simulated behavioural responses (i.e. RTs and choices) of individual participants, with the same number of trials per condition as in the actual experiment. There was a good agreement between the observed data and the model simulations across reward probability levels and choice preferences (Figure 2.3B).

We use Bayesian inference to analyse the posterior distributions of group-level model parameters [27]. To evaluate if a parameter varies substantially between any two conditions, we calculate the proportion of posterior samples in which the parameter value for one condition was greater than the other. To test if a parameter differs from a threshold value, we calculate the proportion of the posteriors greater or smaller than the threshold. To avoid confusion, we use p to refer to classical frequentist p-values, and $P_{p|D}$ to refer to Bayesian inference results based on the proportion of posteriors supporting the testing hypothesis, given the observed data.

For the best fitting model (Figure 2.4A), we compared the posterior estimates of the group-level parameters between conditions (Figure 2.4B and Figure 2.4C). We found strong evidence for choices with high reward probability to have higher mean (v) and standard deviation (S) of the accumulation rate than choices with medium ($v_{high} > v_{medium} : P_{p|D} = 0.999$; $S_{high} > S_{medium} : P_{p|D} = 0.954$) or low medium ($v_{high} > v_{low} : P_{p|D} = 1$; $S_{high} > S_{low} : P_{p|D} > 0.999$) reward probability. The mean and standard deviation of accumulation rates between choices with medium and low reward probabilities were inconclusive ($v_{medium} > v_{low} : P_{p|D} = 0.839$; $S_{medium} > S_{low} : P_{p|D} = 0.877$). Furthermore, there was also strong evidence for a higher mean accumulation rate for the preferred than the non-preferred options ($P_{p|D} = 0.999$), and no evidence for a difference in the standard deviation of the accumulation rate ($P_{p|D} = 0.532$). These results supported the claim that preferred and certain (100%) cues were recalled and

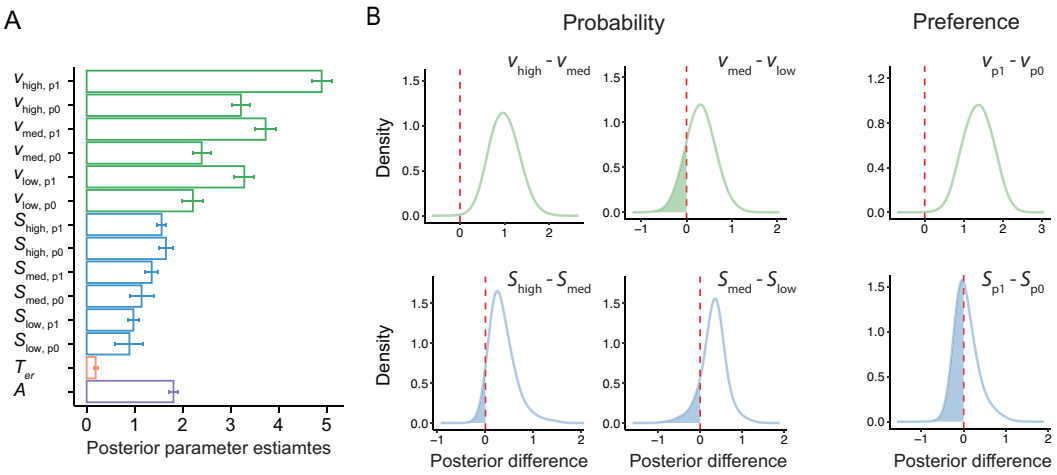

Figure 2.4: *Posterior model parameters and inferences. (A) Group-level LBA model parameters of the best fitting model: means of accumulation rates (v, green), standard deviations of accumulation rates (S - blue), non-decision time (T_{er} - orange) and starting point (A - purple). Error bars represent standard deviations of posterior distributions of parameter values. The means and standard deviations of accumulation rates were shown separately for each reward probability level (high, medium and low) and accumulator (p1 - preferred option; p0 - non-preferred option). (B) Differences of posterior parameter estimates across probability levels (left and middle columns) and preference levels (right column). The proportion of posterior difference distributions above zero suggested higher parameter values for higher probability level or more preferred options.*

processed faster than non-preferred cues. Certain cues were also associated with more variable accumulation rate. Model comparisons further suggested that the latencies of early visual encoding and motor execution were not influenced by reward probability nor preference as the models with varying non-decision time parameter did not fit the data as well.

We also evaluated whether the best fitting model could reproduce qualitative RT patterns in the single-option and unequal choices, which were unseen by the parameter estimation procedure. For unequal trials, the simulated RT showed similar patterns to the observed data, in which choosing between medium and low probability cues led to the longest RT (Figure 2.3C). On the right hand-side of Figure 2.3C, one might notice that high and low reward condition has - according to our model - just intermediate RTs, which is inconsistent with high accuracy of this condition (compare Figure 2.2) by means of the speed-accuracy trade-off. However, both modelling and experimental evidence suggest that the discrimination between certain cues (high reward probability) and the others (medium and low) was much easier task then comparison between the two uncertain ones. Therefore, in the light of the remaining results the model predictions seem to be consistent. For single-option choices, similar to the observed data, higher reward probability and preferred cues were associated with faster RT in simulation. However, simulated RT in single-option choices was longer than the experimental data, suggesting that simple reactions to a single cue may engage distinct cognitive processes beyond the current model.

2.4.3 EEG Results

We focused our EEG analysis on equal trials (with additional control analysis on EEG data from single-option trials), because both reward probability and preference bias played major roles in shaping the behavioural performance of that condition.

Event-Related Potentials

We examine univariate differences in evoked responses between conditions in single EEG electrodes. For each participant, trial-averaged ERPs are calculated from epochs

of equal or single-option choices, with epochs time-locked to reward cue onset. For both equal and single-option conditions, we test for differences in ERPs between three levels of reward probability using a one-way repeated-measures ANOVA. Furthermore, we test for differences in ERPs between preferred and non-preferred choices in equal trials using a paired t-test. We perform statistical tests on all electrodes and all time points. Cluster-based permutation tests (2000 iterations with maximum statistics) are used to correct for multiple comparisons across electrodes and time points [258].

Different reward probability levels produced similar grand-average ERP waveforms during equal (Figure 2.5A) and single-option (Figure 2.5B) choices, with a negative peak in the 100 – 150 ms time window (the N100 component) and a positive peak in the 300 – 400 ms time window (the P300 component).

When assessing the effect of reward probability on ERPs, we found no univariate differences survived the correction for multiple comparisons in equal ($p > 0.552$ at all time points, cluster-level permutation test across electrodes and time points) or single-option trials ($p > 0.175$, cluster-level permutation test). For equal trials, we found no significant difference in ERPs between preferred and non-preferred choices (Figure 2.5C, $p > 0.208$, cluster-level permutation test). Therefore, in the current study, univariate ERPs were not sensitive to reward probability or preferred/non-preferred choices.

Multivariate patterns in equal choices

To decode multivariate information representing reward probability in equal choice trials, we applied the linear SVM on multivariate EEG patterns across all electrodes (see Section 2.3.5). Binary classification between high and medium reward probability was significantly above chance ($p < 0.01$, cluster permutation correction, non-parametric Wilcoxon test) from 144 ms after cue onset (Figure 2.6A). Similarly, the information between high and low reward probability was decodable above chance from 192 ms after cue onset ($p < 0.05$, cluster permutation correction). We found no significant classification accuracy between medium and low reward probability ($p > 0.16$ in all time points, uncorrected). Therefore, choices associated with certain (100%) rewards

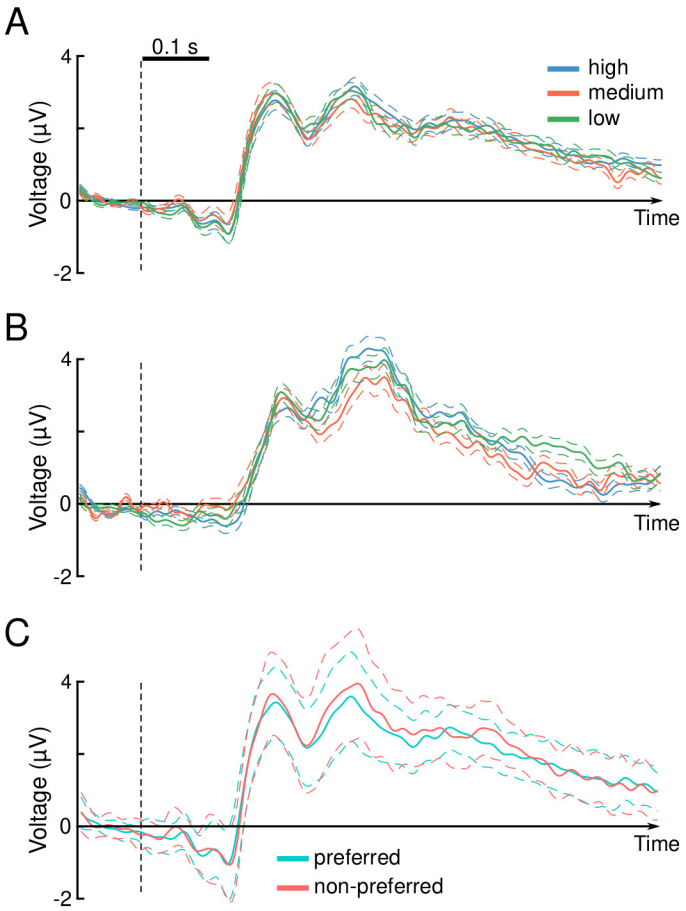

Figure 2.5: *Grand-average stimulus-locked ERPs across all EEG electrodes. (A) ERPs from high (100%), medium (80%) and low (20%) reward probability in equal trials. (B) ERPs from high (100%), medium (80%) and low (20%) reward probability in single-option trials. (C) ERPs from equal trials in which the preferred or non-preferred cue was chosen. In all panels, the dashed lines represent standard errors across participants.*

were distinguishable from those with uncertain reward probabilities.

We applied a similar classification procedure to decode the information between equal trials in which the participants chose their preferred or non-preferred choices across reward probability levels. The information about preferred versus non-preferred choices was decodable from 316 ms to 472 ms after cue onset ($p < 0.009$, cluster permutation correction).

To evaluate the relative importance of each feature (i.e. EEG electrode) to the classification performance, we calculated the weight vector of SVMs. For each classification problem, we retrained the SVM at each time point with all the data included in the training set and obtained the SVM weight vector. The weight vectors were then transformed into interpretable spatial patterns by multiplying the data covariance matrix [168]. The group spatial patterns were calculated by averaging across participants and from all time points which had significant classification accuracy. Relevance spatial patterns based on SVM's weight vector showed that mid-line central and posterior electrodes contained the most information for significant classification (Figure 2.6).

2.4.4 EEG-informed cognitive modelling

P300 component is a strong candidate for a marker of evidence accumulation. Its amplitude has been associated with attention [82], working memory [223] and with task difficulty [223]. Prominent models propose that it reflects build-to-threshold of the decision variable [412, 209] or marks the conclusion of internal decision-making process [284]. Considering that the latency of early visual processing is a part of non-decision time [289], we further hypothesised that the evidence accumulation process initiates at N100 peak latency. This led to a theoretical prediction that the slope of the rise in EEG activity between N100 and P300 peak amplitudes reflected the accumulation rate on a trial-by-trial basis. To validate this prediction, we estimated the N100 and P300 components from single trials of equal choices (Figure 2.7A), using an SVD-based spatial filter to improve the signal-to-noise ratio of single-trial ERPs (see Estimation of single-trial ERP components). This single-trial EEG estimate was then added as a linear regressor (Equation 2.2) of the mean accumulation rate to the

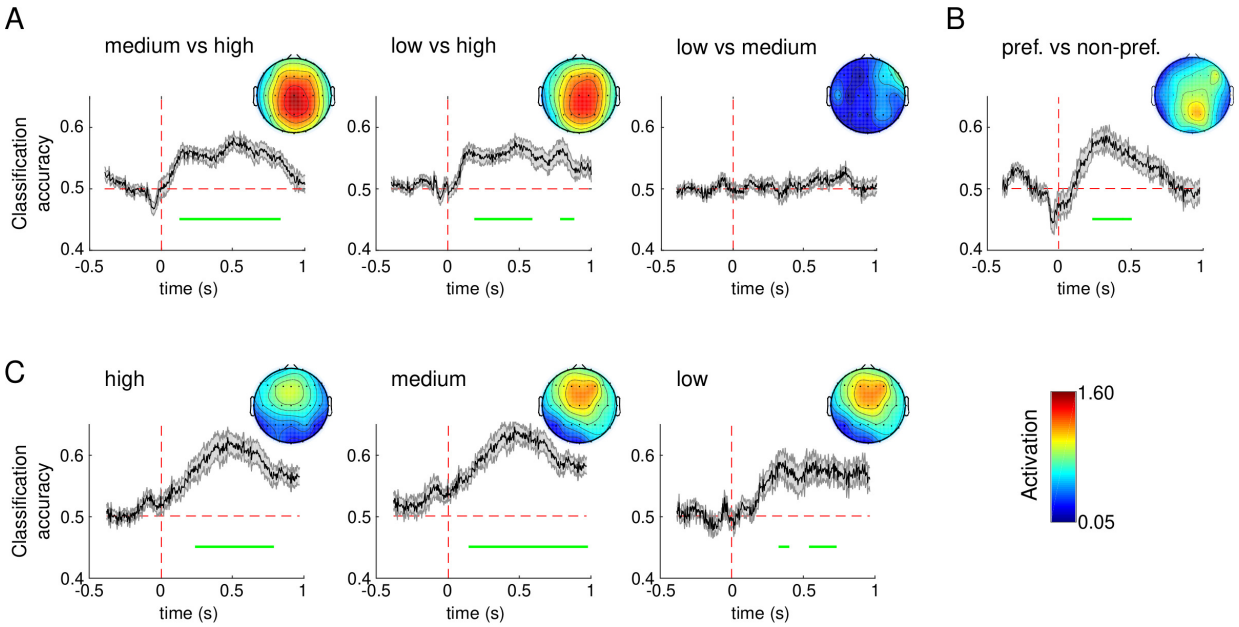

Figure 2.6: *MVPA results.(A) Classification accuracies across time-points between equal choices with different levels of reward probability. (B) Classification accuracies across time-points between equal trials with preferred and non-preferred choices. (C) Classification accuracies across time-points between equal and single-option choices with the same level of reward probability. In all panels, the black lines denote classification accuracies from a stratified 10-fold cross-validation and the gray areas denote standard errors. Significant decoding time windows (green horizontal bars) were determined from cluster-level permutation tests ($p < 0.05$, corrected). Topographic maps represent activation patterns from classification weights, which indicate the contribution of different EEG channels to overall classification accuracies.*

LBA model variant with the best fit to behavioural data.

We used the same MCMC procedure to fit the extended LBA model with the EEG-informed regressor to the equal trial data. The extended LBA model showed good convergence ($\hat{R} \leq 1.02$ for all parameters) and provided a better fit, with a lower LOOIC score 2687 than the model without the EEG-informed regressor (LOOIC score 2796), suggesting that the rising slope of N100-P300 indeed affected the decision process. The posterior estimate of the regression coefficient β provided strong evidence for a positive single-trial effect (Figure 2.7B, $P_{p|D} = 0.983$), indicating that a bigger N100-300 slope is associated with faster accumulation rate.

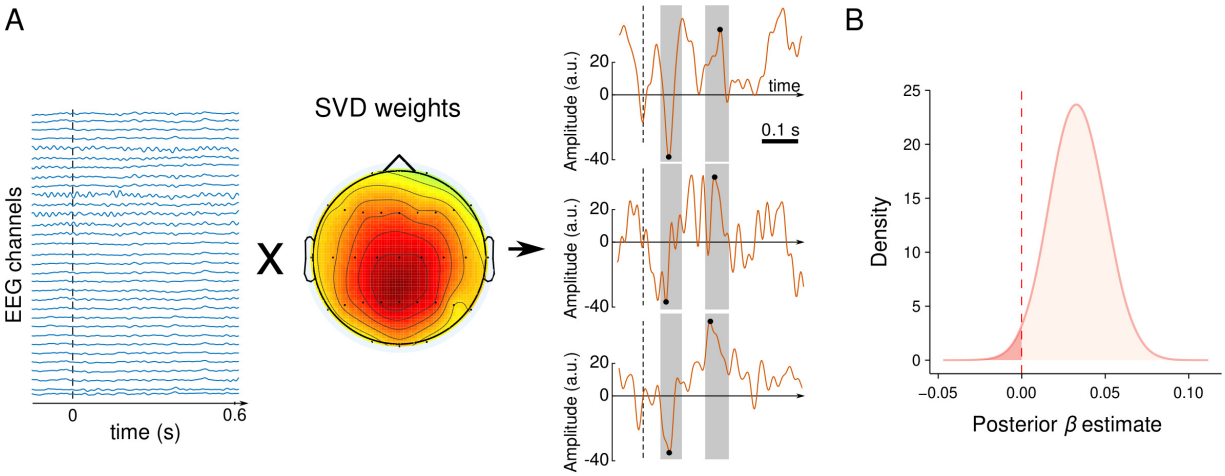

Figure 2.7: *EEG-informed modelling. (A) The schematic diagram of extracting single-trial ERP components. 32-channel EEG signals from a single trial were multiplied by the weights of the first SVD component, calculated from the grand-averaged ERP. Next, the N100 and P300 components in that trial were identified by searching for the peak amplitude in a time of 60-164 ms for the N100 component, and 272-376 ms for the P300 component, respectively. ERP marks in three representative trials were illustrated in the right column of the panel. The ratio between N100-P300 peak amplitude difference and N100-P300 peak latency difference was calculated as a single-trial regressor for modelling. (B) Posterior estimates of the coefficient between the EEG-informed single-trial regressor (i.e. the rising slope of N100-P300 components) and changes in the accumulation rate.*

2.5 Discussion

We provide novel evidence that reward probability and spontaneous preference influence choices between equally probable alternatives and their electrophysiological signatures. We observed two patterns that were consistently distinct at behavioural, cognitive and neural levels: a certainty effect, distinguishing choices between cues with 100% reward probability and cues with uncertain reward probabilities (80% or 20%), and a preference effect, differentiating between equally valued options. At the behavioural level, reward certainty (i.e. 100% reward vs. non 100% rewards) resulted in disproportionately faster reaction times, while preference biased both choice frequency and RT, resulting in more frequent and faster responses for preferred cues. Using hierarchical Bayesian implementation of a cognitive model, we showed that reward certainty and preference bias were associated with changes in the accumulation rate, a model-derived parameter to account for the speed of evidence accumulation during decision-making. At the electrophysiological level, the information of certainty and preference could be reliably decoded from multivariate ERP patterns early during decisions, but not from univariate EEG activities. The accumulation rate was further affected by the slope of the rise in ERPs between the N100 and P300 components on a trial by trial basis. Together, the current study provides insight into neurocognitive mechanisms driving choices in a deadlock situation, where there is no clear advantage in choosing one option over the other.

The certainty effect implies a monotonic but nonlinear relationship between reward probability and RT in equal choices: the difference between certain (100%) and uncertain (80% and 20%) reward was greater than that between the two uncertain conditions. This points to a special status of the 100% reward certainty distinct from lower reward probabilities, as the latter always carries a non-zero risk of no reward. The salient representation of the 100% reward certainty is further highlighted by the lack of significant EEG pattern classification between the two uncertain reward probabilities (i.e. 80% vs. 20%, Figure 2.6A). Here, the certainty effect in rapid voluntary decisions resembles risk-averse behaviour in economic decisions [409], which overweights outcomes with 100% certainty relative to probable ones.

Interestingly, reward probability affected RTs across all trial types. It persisted from equal choices to simple reactions to cue locations in single-option trials. In unequal choices, there was also a negative association between RT and the sum of reward probability of the two choices. Therefore, even though the reward was not contingent upon RT in the current study, we observed a general tendency of accelerating ones' responses in the presence of more certain reward. These results are akin to the effect of reward magnitude, which also demonstrates a facilitating effect on RT [360, 70]. In non-human primates, the phasic activation of dopamine neurons in the ventral midbrain has similar response profiles to changes in reward probability and magnitude [124], suggesting a common mesolimbic dopaminergic pathway underlying different facets of reward processing that affect decision-making.

Bayesian model comparison identified specific effects of reward probability on accumulation rates, highlighting two possible cognitive origins of the certainty effect. First, in equal choices, cues with 100% reward resulted in larger mean accumulation rates than those with uncertain reward probabilities (Figure 2.4A). Accumulation rate has been linked to the allocation of attention on the task [354]. Because reward plays a key role in setting both voluntary (top-down) and stimulus-driven (bottom-up) attentional priority [245, 330, 227, 453], high reward probability may boost the attentional resources allocated to sensory processing for more rapid decisions. Second, reward probability affected the variability of accumulation rates across trials (Figure 2.4A). Higher accumulation rate variability has been associated with better-memorised items [376, 295, 398]. It is possible that stimuli associated with 100% reward were memorised more strongly [268], a hypothesis to be confirmed in future studies.

Furthermore, MVPA of stimulus-locked ERPs showed multivariate EEG patterns distinguishing between cues with 100% reward and other uncertain reward probabilities as early as 150 ms after stimulus onset (Figure 2.7, see also [394]), and model comparisons found no evidence to support for non-decision time to vary between reward probability levels. Considering the average RT of 600 $\sim$ 900 ms in equal choices, our results did not support the latency of post-decision motor preparation, which constitutes a part of the non-decision time [206], to be the source of the certainty effect. This result is consistent with the view that motor action implementation is indepen-

dent of the stimulus value [262]. Instead, the certainty effect possibly originates from evidence accumulation during the decision process, as supported by the changes in the accumulation rate.

When choosing between equally valued options, classical evidence accumulation theories predict a deadlock scenario with a prolonged decision process [49]. This was not supported by recent experimental findings in value-based decisions [307, 391], including the current study, in which equal choices took no longer than unequal ones. Our behavioural, modelling and EEG analyses indicated a preference bias which could effectively serve as a cognitive mechanism to break the decision deadlock. Compared with non-preferred options, preferred decisions facilitated RTs, were associated with larger accumulation rates and evoked distinct EEG multivariate patterns. Here, we did not aim to provide a mechanistic interpretation of preference (i.e. why or how the preference bias originated). Instead, our results demonstrated a consistent presence of preference bias before and after cue-probability re-mapping, independently across reward probabilities and maintained in single-option trials (Figure 2.2C), which we considered as a novel finding in the literature of voluntary choice.

What can induce a preference bias? Because the cue-probability association was initially randomised and later changed within each session, and no differences in shape preference were found, this bias was not due to stimulus salience but established spontaneously [424]. Multiple factors may contribute to the establishment of preferred options. Preference might arise as a function of early choices and outcome frequencies [194, 21], which shape future beliefs or alter the memory trace of certain cue-probability bindings. This interpretation is consistent with an irrationality bias, which favours previously rewarded stimuli, even when controlling for their value [357]. Alternatively, some cue-value associations might be remembered more reliably due to a deliberate cognitive strategy of memory resource allocation.

Our results provide little evidence to either support or refute these hypotheses. However, memory strength alone cannot explain the full set of results in the current study. First, it is worth noting that the stimulus-reward mapping was presented a total of 16 times throughout each session (at the beginning of each block and after every 40 trials), and participants took as much time as they needed before the next set of trials.

Second, the linear mixed-effect models found significant effects of preference on RT only in equal and single-option trials, but not in unequal trials. If we were to believe memorisation of items to be different between two cues of the same reward probability, we would expect this to be reflected also in the unequal condition, which was not the case. Future studies could validate these hypotheses by employing more frequent cue-probability remapping throughout experiments and controlling for memory effects. Furthermore, all trials in the current studies were randomised and participants did not have prior knowledge of upcoming stimuli. One future extension would be to evaluate whether presenting prior information of reward probability in an upcoming trial would modulate boundary separation in voluntary decisions, similar to the effect of prior bias on perceptual decisions [274].

The current study considered a simplified form of decision, in which the amount of reward was fixed (i.e. 10 game points). In traditional value-based decisions assumed by the prospect theory, a decision-maker needs to integrate the value and probability of gain or loss to obtain an expected utility for each option [410]. Together, our results here and previous studies [430] provide converging evidence that both reward value and probability can influence RT in equal choices. This raises the intriguing possibility of our results to be generalized to choices with the same expected utility but the different combinatory of value and probability. Interestingly, the multiattribute extension of the LBA model [408] has been fitted to RTs from such tasks [79], suggesting that our modelling and EEG approaches could also be extended to explore more complex decision problems.

Our study highlights the advantages of EEG-informed cognitive modelling to inform behavioural data. Hierarchical Bayesian parameter estimation of the LBA model provides a robust fit to an individual's behavioural performance with less experimental data needed than other model-fitting methods [419, 448, 472]. By integrating single-trial EEG regressors with the cognitive model, we identified the accumulation rate to be affected by the rate of EEG activity changes between visual N100 and P300 components. This result contributes to a growing literature of EEG markers of evidence accumulation processes, including ERP components [412, 249, 288], readiness potential [252] and oscillatory power [418]. It further consolidates the validity of evidence

accumulation as a common computational mechanism leading to voluntary choices of rewarding stimuli [386, 2, 257], beyond its common applications to perceptually difficult and temporally extended paradigms.

The EEG-informed modelling builds upon the known functional link between the P300 component and evidence accumulation for decisions [311, 421, 412]. A new extension in the current study was to consider the accumulation process begins at the peak latency of the visual N100 component. Theoretically, the delayed initiation of the decision process accounts for information transmission time of $60 \sim 80$ ms from the retina [355]. Single-unit recording concur with this pre-decision delay, as neurons in putative evidence accumulation regions exhibit a transient dip and recovery activity independent of decisions approximately 90 ms after stimulus onset [336]. Practically, our EEG data has a clear N100 component, and time-resolved MVPA identified significant pattern differentiating between task conditions at a similar latency. The relatively early start of the accumulation process in our experiment might be explained by the easily discriminable nature of the cues, consisting of basic shapes with no perceptual noise. Longer visual processing stage has been reported in an experiment involving more complex processing of visual information [289]. Further research could dissect the non-decision time [441, 402] and compare latencies of visual encoding across decision tasks with stimuli at different levels of complexity.

Several issues require further consideration. First, our cognitive modelling was not meant to reproduce all the rich behavioural features in the data. To include sufficient observations for model-fitting, we combined the data before and after cue-probability remapping. As a result, our model did not account for behavioural changes related to cue remapping. Future studies could employ a multi-session design to investigate how learning new cue-probability associations influence model parameters [470].

Second, we focused on the certainty and preference effects by fitting the LBA model only to the data from equal choices. Although simulations indicated that the fitted model provided similar behavioural patterns as in the empirical data in unequal and single-option choices, it was not fitted directly to the experiment data in those two choice conditions. A more parsimonious model for all three types of choices would require additional assumptions, which is beyond the scope of the current study. For

example, to incorporate the large RT discrepancy between equal and single-option choices, one could assume that the urgency signal [46, 395] plays a more dominant role in accelerating RT when no apparent comparisons are needed in single-option choices. This might be a representation of Hick's law - a psychological phenomenon where increasing the number of choices increases the decision time logarithmically [176].

Third, our model selection procedure does not encompass all conceivable model types that might account for this data. Independent accumulation is consistent with findings on brain mechanisms of probability-based choices in humans [224, 357], as well as choice behaviour in rats [292]. Alternative explanations of certainty and preference effects can be provided by urgency gating [396], collapsing threshold [324] or cross-inhibition [298, 414]. Depending on the parametrization, interpretations based on these models could slightly vary. It is also important to note that there may be no straightforward way to disentangle the interpretations provided by these different models [269]. However, more detailed neuroimaging recordings (i.e. with higher temporal or spatial resolution) might help to address this issue [288, 289]. Alternatively, the experimental designs with an additional perturbation push affect the decision time in the perceptual decision-making setting [455, 474]. It has been shown that this can help to distinguish among integrate-to-threshold models [474]. Future extensions of the presented experiment could implement these suggestions to address some of the remaining issues. Nevertheless, these potential limitations, although important, would not challenge the main conclusions of this research.

Finally, our model-based analysis is unavoidably constrained by the choice of model, and one needs to be cautious when extending findings to different models. There is an ongoing debate on how accurately different models can mimic each other when estimating the non-decision time [100, 149, 238]. The DDM, for example, tends to predict longer non-decision times than LBA [105], as well as might be more susceptible to urgency manipulations [119]. Although an extended DDM has been shown to account for magnitude effects [329], the drift rate of a DDM represents the relative signal difference between the two options. As a result, without fitting a new DDM to each condition, the DDM cannot directly describe all conditions in our current study (i.e. the unequal and single-option trials). LBA, on the other hand, assumes an independent accumula-

tor for each option (i.e. each reward cue), offering a parsimonious account to our task and the capacity to produce the qualitative features of responses in all conditions, as demonstrated in our model simulations. It is worth noting that this feature of LBA is shared by other models with multiple accumulators (e.g. the LCA model and the race model). Among those, the LBA model has the simplest form and analytic solutions to the first-passage time, enabling efficient MCMC sampling in Bayesian hierarchical modelling.

In conclusion, when choosing between equally probable reward outcomes, probability and preference selectively modulate the decision processes and their electrophysiological signatures, providing a mechanism for breaking a decision deadlock. These findings extend and substantiate the computational framework of evidence accumulation for voluntary decisions. Our results further highlight the intricate nature of human behaviour, as susceptible to external factors as well as endogenous heuristics.

Chapter 3

Data-driven model: Energy landscape method for MEG data

This chapter is based on the work published in the *Network Neuroscience* [233] in collaboration with Dr Naoki Masuda (conceptualisation, methodology, software, supervision), Dr Khalid Hamandi (data curation, interpretation of the results), Prof. Krish D. Singh (methodology, supervision), Dr Bethany Routley (data curation) and Dr Jiaxiang Zhang (conceptualisation, formal analysis, methodology, project administration and supervision). The author's contribution to this work include conceptualisation, data analysis, methodology, software implementation, data visualisation and writing.

Here, I present a statistical model of interactions between the instantaneous brain states. We use, derived from thermodynamics, pairwise Maximum Entropy Model to determine the occurrence probability of network states in MEG oscillatory power. Energy landscape assigns energy value to each pattern of activations across brain regions. The energy is inversely proportional to the probability of occurrence. We apply this method to the cohort of JME patients and healthy controls matching age and gender. The model provided a good fit to the binarized MEG data in both groups. Patients with JME exhibited fewer local minima of the energy and elevated energy values than controls, predominately in the fronto-parietal network across multiple frequency bands. Furthermore, multivariate features constructed from energy landscapes allowed significant single-patient classification.

3.1 Background

Juvenile myoclonic epilepsy (JME) is the most common syndrome of the wider group of idiopathic generalized epilepsies [451]. Patients with JME often exhibit three main types of seizures: myoclonic, absence and generalized tonic-clonic seizures [452]. Typical JME characteristics are normal or close to normal clinical MRI of the brain and interictal EEG with irregular spike-waves or polyspike-waves with frontal predominance [64]. JME patients are susceptible to seizure precipitation after sleep deprivation, alcohol usage, excise or demanding cognitive processing [90, 458]. JME is a lifelong condition and treatment with antiepileptic drugs is usually necessary.

Although the pathogenetic mechanisms of JME is still not fully understood [36], JME has been recognised as a network disorder affecting brain activity and connectivity that leads to cognitive impairments [72, 452] and personality traits similar to patients with frontal lobe lesions [116]. BOLD functional MRI (fMRI) and diffusion weighted imaging showed hyper-connectivity in the frontal lobe in JME [425, 63]. Electrophysiological data suggests that JME has an impact on multiple functional networks, including the fronto-parietal network (FPN) [452], the default mode network (DMN) [266], and the sensorimotor network (SMN) [76], which may be driven by dysfunctional thalamocortical circuitry [153, 158, 39, 216].

Several sensitive markers from resting EEG and MEG recordings have been identified for classifying patients with epilepsy and predicting seizure onsets, including information entropy [205, 369], Lyapunov exponent [17, 189] and phase plane portraits [189]. These methods describe statistical regularities of electrophysiological signals from a dynamical system perspective, in line with the theoretical account of epileptic seizures as bifurcations from stable states [81]. In JME, however, it is yet unclear whether atypical statistical properties of network activation is present during rest, and if so, whether the changes are frequency specific.

This study addressed these problems by applying a pairwise maximum entropy model (pMEM) approach [462] to source-localised, frequency-specific MEG resting-sate oscillatory activity (Figure 3.2). The pMEM is a statistical model of the occurrence probability of network states, with its parameters being constrained by the network's

Table 3.1: *Demographics of patients with JME and and healthy control participants. (MJ - myoclonic jerks, GTCS - generalised tonic clonic seizures, LEV - leveiracetam, VPA - sodium valproate, LTG - lamotrigine, TPM - topiramate, ZNM - zonisamide.)*

	Patients	**Controls**
Number of participants	26 (8 males)	26 (7 males)
Age median	27	27
Age range	19 - 45	18 - 48
Seizure type (number of patients)	MJ (26) Absences (15) GTCS (26)	-
Anti-epileptic drugs (Number of patients taking the drug)	LEV (13), VPA (12), LTG (5), TPM (4), ZNM (4)	-

regional activity and pairwise regional co-activation from empirical data. According to the principle of maximum entropy, the pMEM is the most parsimonious second-order model of a system with minimum assumptions [197], and it permits multi-stability in a system with meta-stability states [74, 88]. The pMEM has been successfully applied to the collective behaviour of spiking neural networks [400, 356, 390, 40] and BOLD fMRI responses [436, 437, 13, 120]. Here, we extended this theoretical framework to MEG oscillatory activity in three functional networks: FPN, DMN and SMN. Furthermore, based on the fitted pMEM to individual participants, we depicted an energy landscape for each of the networks at theta (4-7 Hz), alpha (8-13 Hz), beta (15-25 Hz) and gamma (30-60 Hz) bands. The energy landscape is a graphical representation of all network states and their energy values [389]. We then compared several quantitative measures obtained from the energy landscapes between JME patients and controls.

3.2 Methods

3.2.1 Participants

Fifty-two subjects participated in the experiment. Demographic and clinical features of the participants are summarized in Table 3.1. Twenty-six patients with JME were

recruited from a specialist clinic for epilepsy at University Hospital of Wales in Cardiff. Consensus clinical diagnostic criteria for JME were used by an experienced neurologist [405]. Inclusion criteria were: (1) seizure onset in late childhood or adolescence with myoclonic jerks, with or without absence seizures, (2) generalised tonic-clonic seizures, (3) normal childhood development as assessed on clinical history and (4) generalised spike wave on EEG and normal structural MRI. Twenty-six healthy control participants with no history of significant neurological or psychiatric disorders were recruited from the regional volunteer panel. All testing was performed with participants' taking their usual medication. The study was approved by the South East Wales NHS ethics committee, Cardiff and Vale Research and Development committees, and Cardiff University School of Psychology Research Ethics Committee. Written informed consent was obtained from all participants.

3.2.2 MEG and MRI data acquisition

All participants underwent separate MEG and MRI sessions. Whole-head MEG recordings were made using a 275-channel CTF radial gradiometer system (CTF Systems, Canada) at a sampling rate of 600 Hz. An additional 29 reference channels were recorded for noise cancellation purposes and the primary sensors were analysed as synthetic third-order gradiometers [427]. Up to three sensors were turned off during recording due to excessive sensor noise. Subjects were instructed to sit comfortably in the MEG chair while their head was supported with a chin rest and with eyes open focus on a red dot on a grey background. For MEG/MRI co-registration, fiduciary markers that are identifiable on the subject's anatomical MRI were placed at fixed distances from three anatomical landmarks (nasion, left and right preauricular) prior to the MEG recording, and their locations were further verified using high-resolution digital photographs. The locations of the fiduciary markers were monitored before and after MEG recording. To ensure that the movement artefacts did not dominate the recording, the average Euclidean distance between fiducials was computed for every participant. There was no significant difference between head movements of the JME and control group ($t(25) = -1.27$, $p = 0.22$) with the mean head shift 0.55 cm. Each

recording session lasted approximately 5 minutes.

Whole-brain T1-weighted MRI data were acquired using a General Electric HDx 3T MRI scanner and a 8-channel receiver head coil (GE Healthcare, Waukesha, WI) at the Cardiff University Brain Research Imaging Centre with an axial 3D fast spoiled gradient recalled sequence (echo time 3 ms; repetition time 8 ms; inversion time 450 ms; flip angle 20°; acquisition matrix 256×192×172; voxel size 1×1×1 mm).

Results presented in this chapter are based upon patient scans. These are confidential and may not be made freely available.

3.2.3 Data pre-processing

Continuous MEG data was first segmented into 2 s epochs. Before segmentation, MEG data was filtered with a 1 Hz high-pass and a 150 Hz low-pass filter to avoid DC step changes between epochs. Every epoch was visually inspected. Those containing major motion, muscle or eye-blink artefact, or interictal spike wave discharges were excluded from subsequent analysis. The artefact-free epochs were then re-concatenated. This artefact rejection procedure resulted in cleaned MEG data with variable lengths between 204 s and 300 s across participants, and the data lengths were comparable between JME patients and controls ($t(50) = 1.38$, $p = 0.17$). The 200 s of cleaned MEG data was used in subsequent analysis. For participants with longer than 200 s cleaned MEG data, a continuous segment of 200 s during the middle of recording session was used.

3.2.4 Source localization of oscillatory activity in resting-state networks

We analysed the MEG oscillatory activity using an established source localisation method for resting-state networks [57, 157, 278].

For each participant, the structural MRI scan was co-registered to MEG sensor space using the locations of the fiducial coils and the CTF software (MRIViewer and MRIConverter). The structural MRI scan was segmented and a volume conduction

model was computed using the semi-realistic model [287].

The preprocessed MEG data was band-passed filtered with a fourth-order zero phase lag Butterworth filter into four frequency bands: theta 4-8 Hz, alpha 8-12 Hz, beta 13-30 Hz, and low-gamma 35-60 Hz [283]. For each frequency band, we downsampled the data to 250 Hz and computed the inverse source reconstruction using an LCMV beamformer on a 6-mm template with a local spheres forward model in Fieldtrip (version 20161101, http://www.fieldtriptoolbox.org).

The atlas-based source reconstruction was used to derive virtual sensors for every voxel in each of the 90 regions of the Automated Anatomical Label (AAL) atlas [180]. Each virtual sensor's time course was then reconstructed.

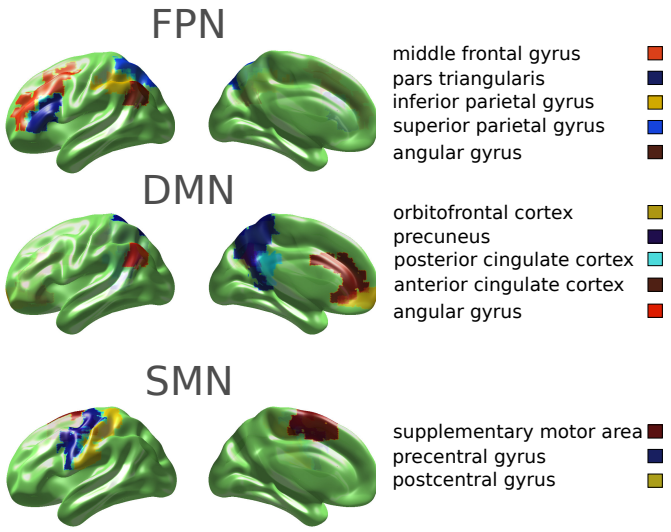

Figure 3.1: *The regions of interest (ROIs) of three resting state networks: the frontoparietal network (FPN), the default mode network (DMN) and the sensori-motor network (SMN). The ROIs were obtained from the 90 AAL atlas [180].*

We focused our analysis on three resting-state networks (Figure 3.1): the frontoparietal network (FPN), the default mode network (DMN), and the sensorimotor network (SMN) in which electrophysiological changes had been reported in patients with epilepsy [266, 76, 452]. Each resting-state network comprised of bilateral regions of interest (ROIs) from the AAL atlas identified in previous studies [392, 339]. The FPN included 10 ROIs: middle frontal gyrus (MFG), pars triangularis (PTr), inferior

parietal gyrus (IPG), superior parietal gyrus (SPG) and angular gyrus (AG). The DMN included 10 ROIs: orbitofrontal cortex (OFC), anterior cingulate cortex (ACC), posterior cingulate cortex (PCC), precuneus (pCUN) and AG. The SMN included 6 ROIs: precentral gyrus (preCG), postcentral gyrus (postCG) and supplementary motor area (SMA). For each ROI, its representative time course was obtained from the voxel in that ROI with the highest temporal standard-deviation. The mean MEG activities of the ROIs of each network were not significantly different between JME patients and controls (FPN: $F(1,50) = 0.75, p = 0.39$; DMN: $F(1,50) = 0.21, p = 0.65$; SMN: $F(1,50) = 0.15, p = 0.70$).

To calculate the oscillatory activity, we applied Hilbert transformation to each ROI's time course, and computed the absolute value of the analytical representation of the signal to generate an amplitude envelope of the oscillatory signals in each frequency band.

3.2.5 Pairwise maximum entropy model of MEG oscillatory activity

During rest, different brain regions exhibit pairwise co-occurrence of oscillatory activity [185] and rapid changes of brain network states [372]. To obtain an estimate of network state transitions and their probabilities, we fitted a pMEM to individual participant's MEG data, separately for each resting-state network and each frequency band.

According to the principle of maximum entropy, among all probabilistic models describing empirical data, one should choose the one with the largest uncertainty (i.e. entropy), because it makes the minimum assumptions of additional information that would otherwise lower the uncertainty [462]. The pMEM estimates the state probability of a network, with its regional activity and regional co-occurrence to be constrained by empirical data. It is equivalent to the Ising model in statistical mechanics [40]. In neuroscience, the pMEM was firstly used in a seminal study for fitting the distribution of neuronal spiking activity across cells [356, 400]. More recently, the same method was used in fMRI study where it was shown that the model is capable of estimating the underlying functional connectivity with a higher accuracy than other functional

connectivity methods [436]. Later studies using the pMEM for fMRI have identified key characteristics of brain state transition [204], perceptual metastability [438], and the effects of ageing [120]. A further advantage of using the pMEM is that various statistical physics theory of the model is available, potentially contributing to the understanding of multivariate data when they are fitted with the pMEM [40]. The current study used this approach to unveil differences between the JME patients and controls in large-scale brain networks (Figure 3.2). Below we outlined the theoretical background and the fitting procedure. A more detailed description of the pMEM modelling and subsequent energy landscape analysis is available elsewhere [389].

Consider a resting-state network consisting of N ROIs. For each real-valued ROI's signal, we thresholded the ROI's Hilbert envelope according to the median of the amplitude. Data points above the threshold were denoted as high oscillatory power $(+1)$, and data points below the threshold were denoted as low oscillatory power (-1). The oscillatory activity in ROI i $(i = 1, ..., N)$ at time t was transformed to a binary time series $r_i(t)$, with $r_i(t) = +1$ for high oscillatory activity and $r_i(t) = -1$ for low oscillatory activity. The activity pattern of a N-dimensional binary vector $\mathbf{s}(t) = [r_1(t), r_2(t), ..., r_N(t)]$, representing the state of the network at time t.

The N-ROI network has a total of 2^N possible states $\mathbf{s}_k$ $(k = 1, ..., 2^N)$. From the binarized oscillatory activity, we calculated the probability of occurrence of each network state, denoted by $P_{\text{emp}}(\mathbf{s}_k)$. We further calculated the empirical average activation rate for each ROI $\langle r_i \rangle_{\text{emp}}$ and the pairwise co-occurrence between any two ROIs $\langle r_i r_j \rangle_{\text{emp}}$:

$$\langle r_i \rangle_{\text{emp}} = \frac{1}{T} \sum_{t=1}^{T} r_i(t), \tag{3.1}$$

$$\langle r_i r_j \rangle_{\text{emp}} = \frac{1}{T} \sum_{t=1}^{T} r_i(t) r_j(t), \tag{3.2}$$

where T denotes the number of timepoints in the data. The fitting procedure aimed to identify a pMEM model that preserves the constraints in Equations (3.1) and (3.2) and reproduces the empirical state probability $P_{\text{emp}}(\mathbf{s}_k)$ with the maximum entropy. It

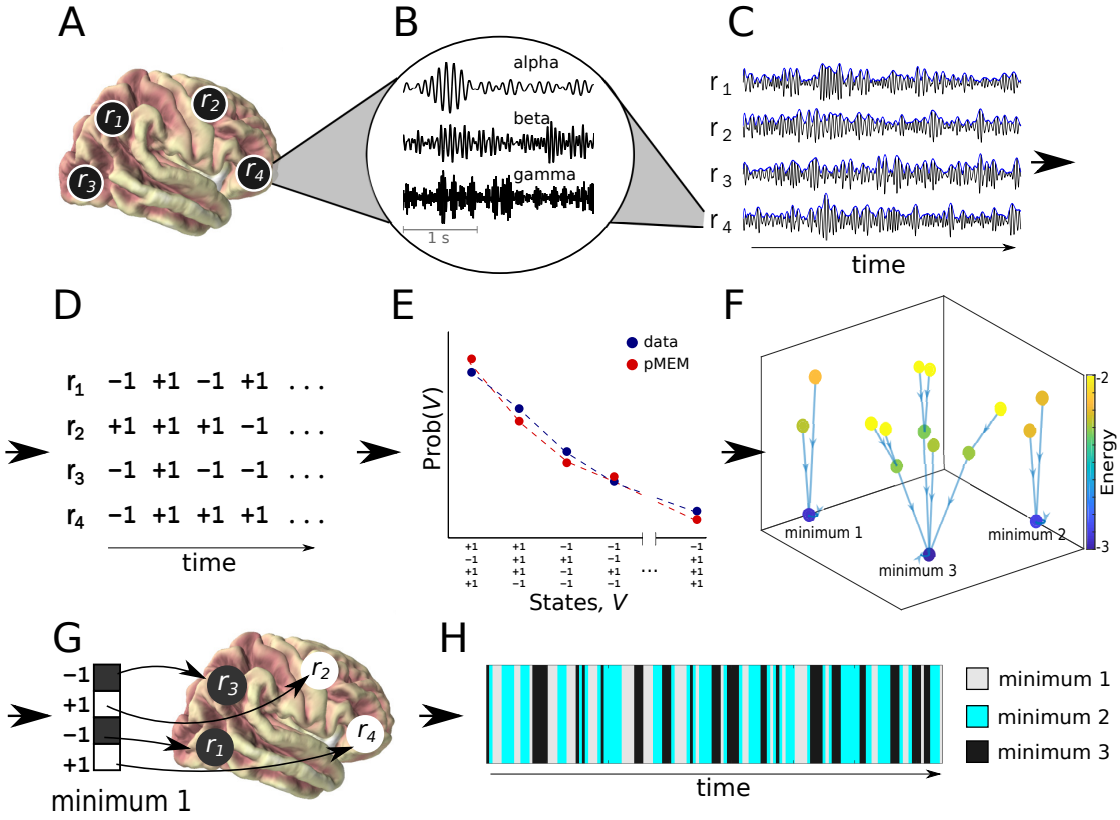

Figure 3.2: *Illustration of the energy landscape analysis on a network of 4 ROIs. (A) Selection of ROIs from the source-space signals. (B) Signal filtering in frequency bands of interest. (C) Envelope extraction using the absolute value of the analytical representation of the signal. (D) Binarization of the data; (E) Fitting the pMEM to match the empirical data distribution of binarized network states. (F) Determining the relationships between network states using the Dijkstra algorithm on energy values. (G) Interpretation of local minima of the energy on the anatomical level. (H) Simulation of the occurrence of network states belonging to different basins.*

is known that the pMEM follows the Boltzman distribution [462], given by

$$P_{\text{pMEM}}(\mathbf{s}_k|\mathbf{h},\mathbf{J}) = \frac{\exp\left(-E(\mathbf{s}_k)\right)}{\sum_{k'=1}^{2^N}\exp\left(-E(\mathbf{s}_{k'})\right)}, \tag{3.3}$$

where $E(\mathbf{s}_k)$ represents the energy of the network state $\mathbf{s}_k$, defined by

$$E(\mathbf{s}_k) = -\sum_{i=1}^{N}\mathbf{h}_i r_i(\mathbf{s}_k) - \frac{1}{2}\sum_{i=1}^{N}\sum_{\substack{j=1\\ j\neq i}}^{N}\mathbf{J}_{ij} r_i(\mathbf{s}_k) r_j(\mathbf{s}_k), \tag{3.4}$$

and $r_i(\mathbf{s}_k)$ refers to the ith element of the network state $\mathbf{s}_k$. $\mathbf{h}$ and $\mathbf{J}$ are the model parameters to be estimated from the data: $\mathbf{h} = [h_1, h_2, ..., h_N]$ represents the bias in the intensity of the oscillatory activity in each ROI; $\mathbf{J} = [J_{11}, J_{12}, ..., J_{NN}]$ represents the coupling strength between two ROIs. The average of the activation rate $\langle r_i\rangle_{\text{mod}}$ and pairwise co-occurrence $\langle r_i r_j\rangle_{\text{mod}}$ expected by the pMEM are given by:

$$\langle r_i\rangle_{\text{mod}} = \sum_{k=1}^{2^N} r_i(\mathbf{s}_k) P_{\text{pMEM}}(\mathbf{s}_k|\mathbf{h},\mathbf{J}), \tag{3.5}$$

$$\langle r_i r_j\rangle_{\text{mod}} = \sum_{k=1}^{2^N} r_i(\mathbf{s}_k) r_j(\mathbf{s}_k) P_{\text{pMEM}}(\mathbf{s}_k|\mathbf{h},\mathbf{J}). \tag{3.6}$$

We used a gradient ascent algorithm to iteratively update $\mathbf{h}$ and $\mathbf{J}$, until $\langle r_i\rangle_{\text{mod}}$ and $\langle r_i r_j\rangle_{\text{mod}}$ match $\langle r_i\rangle_{\text{emp}}$ and $\langle r_i r_j\rangle_{\text{emp}}$ from the observed data, with a stop criterion of 5×10^6 steps. In each iteration, the updates of the parameters were given by $\mathbf{h}_i^{\text{new}} = \mathbf{h}_i^{\text{old}} + \epsilon(\langle r_i\rangle_{\text{emp}} - \langle r_i\rangle_{\text{mod}})$ and $\mathbf{J}_{ij}^{\text{new}} = \mathbf{J}_{ij}^{\text{old}} + \epsilon(\langle r_i r_j\rangle_{\text{emp}} - \langle r_i r_j\rangle_{\text{mod}})$. The learning rate ϵ was set to 10^{-8}.

As in previous studies [437, 389, 120], we used an accuracy index:

$$d = (D_1 - D_2)/D_1 \tag{3.7}$$

to quantify the goodness of fit of the pMEM (Figure 3.3), where

$$D_2 = \sum_{k=1}^{2^N} P_{\text{emp}}(\mathbf{s}_k) \log_2(P_{\text{emp}}(\mathbf{s}_k)/P_{\text{pMEM}}(\mathbf{s}_k)) \tag{3.8}$$

is the Kullback-Leibler divergence between the probability distribution of the pMEM and the empirical distribution of the network state. D_1 represents the Kullback-Leibler divergence between the independent MEM and data. By definition, the independent MEM is restricted to have no pairwise interaction (i.e. $\mathbf{J} = 0$). Therefore, d represents the surplus of the fit of the pMEM over the fit of the independent model. The index $d = 1$ when the pMEM reproduces the empirical distribution of activity patterns and interactions without errors, and $d = 0$ when the pairwise interactions do not contribute to the description of the empirical distribution.

3.2.6 Energy landscape of resting-state network dynamics

The pMEM parameters $\mathbf{h}$ and $\mathbf{J}$ determine the energy $E(\mathbf{s}_k)$ of each network state $\mathbf{s}_k$ ($k = 1, ..., 2^N$), given by Equation (3.4). It is worth noting that, the current study used pMEM as a statistical model to be constructed from the MEG data, not as its literal notion from statistical physics. We did not claim that $E(\mathbf{s}_k)$ represents the metabolic or physical energy of a biological system. Instead, the concept of the energy of a resting-state network stems from the information theory [389]. Here, $E(\mathbf{s}_k)$ indicates the model prediction of the inverse appearance probability of the state $\mathbf{s}_k$ under the empirical constraints of regional activity (parameter $\mathbf{h}$) and regional interactions (parameter $\mathbf{J}$). For instance, if $E(\mathbf{s}_i) < E(\mathbf{s}_j)$, the pMEM predicts that the network activity pattern is more likely to be at the state $\mathbf{s}_i$ than $\mathbf{s}_j$.

For each resting-state network and each frequency band, we depicted an energy landscape as a graph of the energy function across the 2^N possible network states $\mathbf{s}_k$, characterising state probabilities and state transitions from the perspective of attractor dynamics [437]. Because the computational cost increases dramatically with the size of a network, we estimated an energy landscape separately for each resting-state network.

The energy landscape of a network was defined by two factors: the energy $E(\mathbf{s}_k)$

of each network state, and an adjacency matrix defining the connectivity between network states. Two states were defined to be adjacent, or directly connected, if and only if just one ROI of the network had different binarized oscillatory activity (high vs. low). In other words, two states are adjacent when they have a Hamming distance of 1 between their binary activity vectors. For example, for a network with 4 ROIs, states $[-1, -1, -1, +1]$ and $[-1, -1, +1, +1]$ are adjacent, and states $[-1, -1, -1, +1]$ and $[-1, -1, +1, -1]$ are not.

3.2.7 Quantitative measures of energy landscape

We used three measures to understand the differences in the energy landscape between JME patients and healthy controls: (1) the number of local energy minima of within network dynamics, (2) the relative energy of the local minima, and (3) the generative basin duration at significant minima.

Number of local energy minima

A local energy minimum was defined as the network state with a lower energy value than all its adjacent neighbouring states. Because lower energy corresponds to a higher probability of occurrence, network states of local minima can be likened as attractors in attractor dynamics. For each participant, we exhaustively searched through the 2^N network states to identify all the local minima of the participant's energy landscape. We then compared the number of local energy minima between JME and control groups (Figure 3.4).

Relative energy of the local minima

The number of local minima is determined by the energy difference between network states and their adjacent neighbours (i.e. a minimum has a lower energy level than all its neighbours). On the other hand, the energy value of a specific state is determined by its occurrence probability (Equation 3.3). Therefore, theoretically, the two measures had no direct dependency. The energy values of local minima on aggregated energy landscapes indicate the ease of transition from one stable state to another [389, 204].

We calculated the mean energy $\overline{E}(\mathbf{s}_k)$ of each network state $\mathbf{s}_k$ averaged across all participants. Then, we used the mean energy to depict an aggregated landscape, which allowed us to identify common energy minima shared between JME patient and control groups. To test whether each local minimum in the aggregated energy landscape is a characteristic feature of the observed data, we conducted non-parametric permutation tests on the mean energy values. For each resting-state network and each frequency band, we conducted 1000 permutations. In each permutation, we randomly shuffle the pMEM parameters $\mathbf{h}$ and $\mathbf{J}$ (between ROIs and ROI pairs, respectively) that were fitted to individual participants. We then calculated an averaged energy landscape across all participants based on the shuffled parameters. This gave us a sampling distribution of the energy of each network state, under the null hypothesis that the energy values are not related to the observed oscillatory activities or observed pairwise regional co-occurrence. For each local minimum of the aggregated landscape, the level of significance (p-value) of that local minimum's energy was estimated by the fraction of the permutation samples that were higher than the mean energy $\overline{E}(\mathbf{s}_k)$ of that network state in the empirical data without shuffling. To account for the multiple statistical tests that were performed for all the local minima of each network, we evaluated the results using a Bonferroni-corrected threshold ($p < 0.05$) for significance.

Because the shape of an energy landscape was partly determined by the global minimum [437, 120], for each participant, we calculated the energy difference between a significant local minimum and the global energy minimum (i.e. the state with the lowest energy value on the landscape). We then compared this relative energy of the within-network local minima between JME patients and healthy controls. From the networks with significant alternations of relative energy values in JME patients, we constructed a disconnectivity graph to describe clusters of local minima and the relationships between them [30], where a cluster represents a group of local minima with high probabilities of subsequent occurrences [389].

Basin duration at significant minima

On the aggregated energy landscape, the energy basin for each significant local minimum was identified using an existing method [437]. We started at an arbitrary network state and moved downhill on the energy landscape to one of its neighbouring state with the lowest energy, until a local minimum was reached. The starting state is then assigned to the basin of the resulting local minimum. We repeated this procedure for all network states as the starting state.

We used the fitted pMEM as a generative model to simulate the dynamical changes in each resting-state network, and estimated the duration of the basin of each local minimum in the simulated dynamics. Similar to previous studies, we employed the Metropolis-Hastings algorithm to simulate time courses of network activity [166]. Each simulation started with a random network state $\mathbf{s}_k$. On each time step, one of the current state's N neighbouring state $\mathbf{s}_{k'}$ was selected with a probability of $1/N$ as the potential target of state transition, and the state transition occurred with a probability of $\exp\left[E(\mathbf{s}_k) - E(\mathbf{s}_{k'})\right]$ when $E(\mathbf{s}_{k'}) > E(\mathbf{s}_k)$ or 1 otherwise. For each participant, each network, and each frequency band, we simulated 20,000 time steps, and discarded the first 1000 time steps to minimise the effect of initial condition. From the remaining 19000 time steps, we calculated the proportion of duration of the network states that belongs to each energy basin.

3.2.8 Classification of individual patients based on energy values

To investigate the predictive power of pMEM energy measures, we used a support vector machine (SVM) classifier with a radial basis function (RBF) kernel and a leave-one-out cross-validation procedure to classify individual JME patients and controls. The trade-off between errors of the SVM on training data and margin maximization was set to 1. For each resting-state network and each frequency band, the feature space for classification included the energy values of all the significant local minima. In each cross-validation fold, one participant was first removed and the remaining participants' data were used as a training set to train the classifier. To avoid over-fitting,

the feature space (i.e. the local minima) was identified from the aggregated energy landscape constructed from the participants in the training set. The participant left out was then classified into one of the two groups (patients or controls). Classification performance was evaluated by the proportion of correctly classified participants over all cross-validations.

We used permutation tests to evaluate the classification results. The significance of each classification was determined by comparing the observed classification accuracy with its null distribution under the assumption of no difference between patients and controls. The null distribution was generated by 1000 random permutations of leave-one-out classification results, with group labels shuffled in each permutation. We obtained a permutation p-value by calculating the fraction of the permuted samples exceeding the observed classification accuracy.

3.3 Analysis

A summary of participant demographics and clinical characteristics is given in Table 3.1. The JME and control groups were well matched for age ($F(1, 51) = 0.13$, $p = 0.72$) and gender ($p = 0.31$, χ^2 test). For each participant, we performed source localization of pre-processed MEG resting-state data and estimated oscillatory activity (i.e. Hilbert envelope) in each of the 90 ROIs from the AAL atlas, separately in the theta (4-8 Hz), alpha (8-12 Hz), beta (13-30 Hz), and low-gamma (35-60 Hz) bands. We focused our analysis on the differences between JME patients and controls in three resting-state networks (Figure 3.1): the fronto-parietal network (FPN), the default mode network (DMN), and the sensorimotor network (SMN).

3.3.1 Fitting of pairwise maximum entropy models (pMEM) to MEG oscillatory activity

We thresholded an ROI's oscillatory amplitude at each time point t to assign the binary states of 'high' $(+1)$ or 'low' (-1) activity. The state of a network at time t was then represented by a binary vector, consisting of the binarized activity of all the ROIs in

the network. We fitted a pMEM to the series of binarized network oscillatory activities, separately for each participant, each resting-state network, and each frequency band (Equation 3.3, and see Section 3.2 for details). For a network of N ROIs, there are a total of 2^N possible states. The pMEM provides a statistical model of the occurrence probabilities of the 2^N network states, while it satisfies the empirical constraints of mean regional activities at each ROI and pairwise co-occurrence between each pair of ROIs within the network.

To evaluate the model fit, we compared the predicted and observed occurrence probabilities of the 2^N possible network state, averaged across the participants in each group. There was a good agreement between the model predictions and observed data across networks in the JME ($R^2 > 0.90$ in all networks and frequency bands, based on a log-log regression, Figure 3.3) and control groups ($R^2 > 0.89$). We further used an accuracy index to quantify the goodness of fit of the pMEM (Equation (3.7)). The accuracy index was calculated as the percentage of improvement of the pMEM fit to the empirical data compared with a null model, which assumed no pairwise co-occurrence between ROIs (i.e. an independent maximum entropy model). The pMEM achieved high accuracy indices in both JME patients and controls (Figure 3.3). A Mann-Whitney U-test on accuracy indices showed no significant main effect of group (JME vs. controls: $U = 266.0$, $p = 0.19$), suggesting the robustness of the pMEM on MEG oscillatory activities in both patients and controls. As determined by non-parametric repeated measures Friedman test, there were main effects of the networks ($\chi^2 = 87.5$, $p < 0.00001$) and the frequency bands ($\chi^2 = 46.27$, $p < 0.00001$), suggesting that the distinct properties of the networks and information carried by the frequency bands affected the goodness of fit.

3.3.2 Inferences from pMEM energy landscape

Fitted pMEM yielded an energy value for each network state (Equation 3.4). We used energy values from the pMEM to depict an energy landscape of the network. The energy landscape is a graph representation of energy values from all possible network states (Figure 3.2F). We defined two network states being adjacent if there is one and

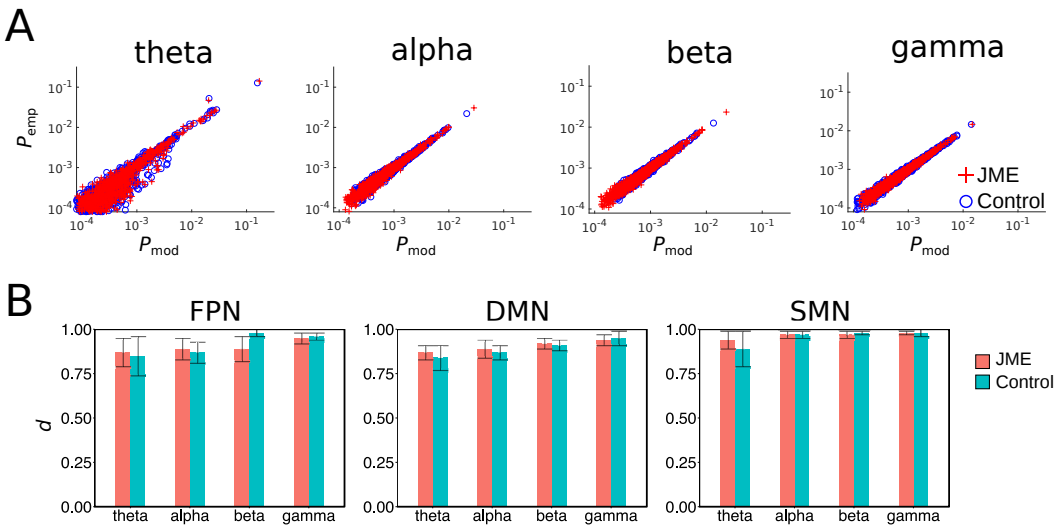

Figure 3.3: *The pMEM fitting. (A) The occurrence probability of each network state of the FPN from the fitted pMEM (P_{mod}) was plotted against that from the empirical data (P_{emp}). Each data point was averaged across JME patients (red) and controls (blue). (B) The averaged accuracy index d in the JME and control groups for each network and frequency band. Error bars denote the standard errors across participants.*

only one ROI whose binarized activity (i.e. $+1$ or -1) is the opposite between the two states. According to the pMEM (Equations (3.3) and (3.4)), network states with a higher energy would occur less frequently than those with a lower energy. As a result, transitions from high to low energy states would more likely to occur than that from low to high states. Here, we examined the differences in three quantitative measures of energy landscape between patients with JME and controls: (1) the number of energy minima, (2) the relative energy values at the local minima, and (3) the generative basin duration at significant minima.

Number of energy minima

We located local minima on the energy landscape, defined as the network states with lower energy than all their adjacent states. Because a local minimum state would have a higher occurrence probability than all of its neighbouring states, transitions of network states near an energy minimum is akin to a fixed point attractor in a deterministic dynamical system, and the number of energy minima quantifies the degree

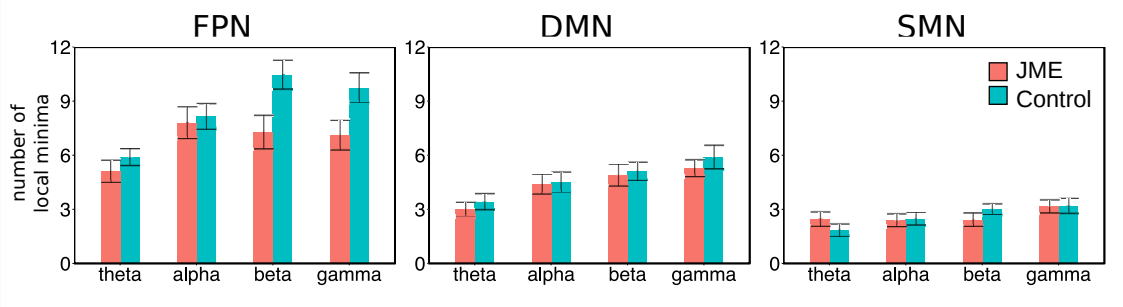

Figure 3.4: *The averaged number of local minima in the JME and control groups. Error bars denote the standard errors across participants.*

of multi-stability of a network.

We calculated the number of local minima for each participant (Figure 3.4) and compared it between groups, resting-state networks, and frequency bands with a repeated-measures ANOVA. Compared with controls, JME patients had significantly less local energy minima ($F(1, 50) = 7.602, p = 0.008$). Across all participants, there were significant main effects of the resting-state network ($F(1.52, 76.25) = 99.89, p < 0.00001$, Greenhouse corrected) and frequency band ($F(2.83, 141.57) = 21.08, p < 0.00001$). No significant network by group ($F(1.52, 76.25) = 3.15, p = 0.07$) or frequency band by group ($F(2.83, 141.57) = 2.12, p = 0.11$) interaction was observed. These results suggested that MEG oscillatory activities in JME patients had altered multi-stability in some networks and frequency bands. The numbers of local minima in beta and gamma bands were more substantially elevated in FPN. The dependence between frequency band and shape of the energy landscape needs to be further investigated.

Relative energy values of the local minima

To identify common energy minima at the group level, we averaged across all participants the energy value of each network state and identified the energy minima on the aggregated energy landscape. In all the three resting-state networks (FPN, DMN and SMN) and all frequency bands, permutation tests showed that the energy values of two network states, 'all off' (i.e. all ROIs had low oscillatory activities $[-1, -1, ..., -1]$) and 'all on' (i.e. all ROIs had high oscillatory activities $[+1, +1, ..., +1]$), did not differ significantly from those from a randomly shuffled energy landscape ($p > 0.88$,

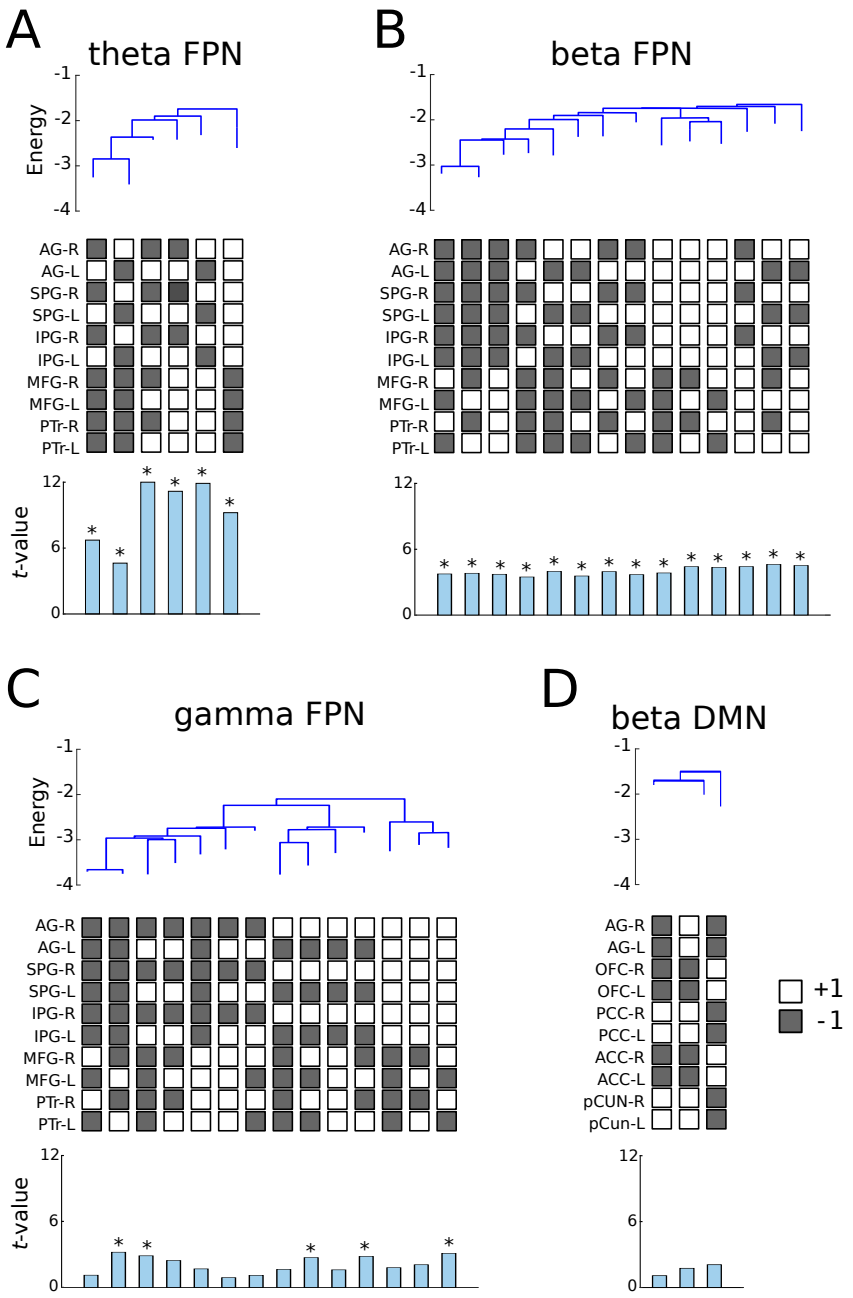

Figure 3.5: *Relative energy values of the local minima in (A) theta FPN; (B) beta FPN; (C) gamma FPN and (D) beta DMN. At the top of each panel, the so-called* ***disconnectivity graph*** *shows the relative energy values of local minima from the aggregated energy landscape across all participants. The end of each branch on the disconnectivity graph represents a local minimum. The middle of each panel shows the network states of the corresponding local minima. White boxes denote high oscillatory activity (i.e. a binary value of* $+1$*) and grey box denote low oscillatory activity (i.e. a binary value of* -1*). The bottom of each panel shows the t-values from two sample t-tests (JME patients vs. controls) on the relative energy values of each local minimum. Asterisks indicate significant difference between JME patient and control groups (*$p < 0.05$*, FDR corrected).*

Bonferroni corrected). That is, the observed energy values at these two minima were not significantly sensitive to regional activation and pairwise coactivation in empirical data (see Section 3.2 for details). In addition, the 'all off' state was also the global minimum of the energy landscape at both group and individual levels, which had the lowest energy value in all network states.

For each significant local minimum state that survived the permutation test, we calculated the relative energy difference between the local minimum and the 'all off' state (i.e. the global minimum) for the individual participants. Then, we compared the obtained relative energy values between the JME and control groups. This subtraction step controlled for the individual variability in the occurrence probability of the global minimum state [437].

In the FPN, the relative energy values at the local minima were significantly higher in JME patients than in controls in the theta-band (Figure 3.5A, $F(1,50) = 18.90, p < 0.0001$), beta-band (Figure 3.5B, $F(1,50) = 15.43$, $p = 0.0002$), and gamma band (Figure 3.5C, $F(1,50) = 7.2558$, $p = 0.009$), but not in the alpha band ($F(1,50) = 0.80$, $p = 0.37$). The aggregated energy landscapes in the beta and gamma bands contained the same set of 14 local minima. Post-hoc tests showed that all the 14 local minima states had higher relative energy values in JME patients than controls in the beta band, and 5 of the 14 local minima states showed significant group differences in the gamma band ($p < 0.05$, Šidák correction). The theta-band energy landscape contained 6 local minima states, which were a subset of the 14 local minima in the higher frequency bands, and all had higher relative energy values in JME patients than controls.

In the DMN, there were trends of higher relative energy values in the JME patients than controls in the beta-band ($F(1,50) = 3.68$, $p = 0.06$) and gamma-band ($F(1,50) = 3.81$, $p = 0.06$), and no significant difference in the theta-band ($F(1,50) = 0.01$, $p = 0.92$) or alpha-band ($F(1,50) = 0.82$, $p = 0.37$). One local-minima in the beta-band, comprised of co-activation in bilateral mPFC and ACC (Figure 3.5D), showed a group difference in post-hoc tests at an uncorrected threshold ($t(50) = 2.34$, $p = 0.03$). In the SMN, there was no significant group difference in the relative energy values (theta: $F(1,50) = 1.26$, $p = 0.27$; alpha: $F(1,50) = 0.06$, $p = 0.81$; beta: $F(1,50) = 0.002$, $p = 0.97$; gamma: $F(1,50) = 0.12$, $p = 0.73$).

Overall, JME patients had higher relative energy values than controls in selective resting-state networks and frequency bands. This result indicates that some local minima on the aggregated energy landscape were less stable (i.e. having a higher energy level) in JME patients than controls.

Basin duration at significant minima

Each local minimum of an energy landscape is accompanied by a basin, which includes the local minimum itself and its neighbouring states from which the local minimum is relatively easily reached [389]. Therefore, the proportion of time for which each basin is visited gives a granular description of network dynamics. For each of the group-level significant minima on the the aggregated energy landscape, we identified all the network states belonging to the same basin. For each participant, we then used the fitted pMEM to numerically simulate network dynamics, and calculated the proportion of time for which the simulated network activities visit each basin. The rationale to simulate basin durations is twofold. First, our simulation demonstrated the feasibility of the derived energy landscape to be used as a generative model of network dynamics. Second, because we removed MEG epochs strongly affected by artefacts, the source reconstructed data was not fully continuous in time, and hence basin duration estimated directly from the empirical data would be less accurate.

In the FPN, simulations showed that network dynamics in JME patients contained shorter basin duration at those significant local minima than controls in the theta ($F(1,50) = 42.72$, $p < 0.000001$), beta ($F(1,50) = 10.49$, $p = 0.002$) and gamma ($F(1,50) = 6.18$, $p = 0.016$) bands, but not in the alpha band ($F(1,50) = 3.92$, $p = 0.053$). There was no significant group difference in the basin duration in the DMN (theta: $F(1,50) = 0.015$, $p = 0.90$; alpha: $F(1,50) = 2.67$, $p = 0.11$; beta: $F(1,50) = 2.76$, $p = 0.10$; gamma: $F(1,50) = 3.12$, $p = 0.08$) or SMN (theta: $F(1,50) = 0.09$, $p = 0.76$; alpha: $F(1,50) = 0.31$, $p = 0.58$; beta: $F(1,50) = 1.59$, $p = 0.21$; gamma: $F(1,50) = 1.25$, $p = 0.27$).

3.3.3 Classification of patients

We used a leave-one-out cross validation procedure for a binary classification of participant groups (JME patients and healthy controls), using the relative energy values of local minima as features. Consistent with the group comparisons (Figure 3.5), the relative energy values obtained from the fitted pMEM showed significant predictive power, with high classification accuracies from theta-band FPN (92.3%, $p < 0.001$, permutation test) and gamma-band FPN (67.3%, $p = 0.012$) (Figure 3.6). The classification based on the energy values from theta-band FPN achieved high specificity (89.3%) and sensitivity (94.8%). For the classification based on gamma-band FPN features, the specificity and sensitivity were 71.4% and 64.5%, respectively. The classification accuracy in the SMN, DMN and other frequency bands of the FPN was not significant ($p > 0.26$, permutation test).

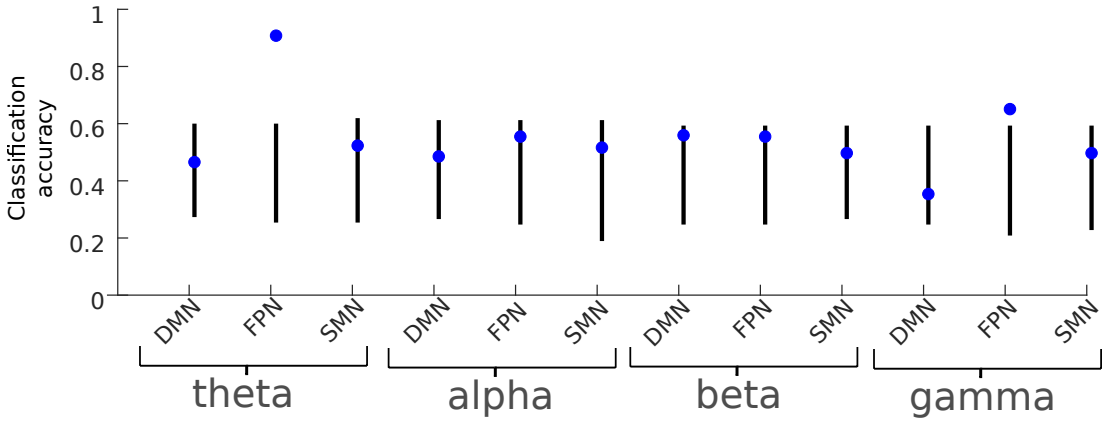

Figure 3.6: *SVM leave-one-out binary classification accuracy of JME patients versus controls. The energy values of the local minima were used as features for classifiers. Blue data points denote the mean classification accuracy. Black lines denote the 95% confidence level under the null hypothesis of no difference between the groups, based on 1000 permutations of randomly shuffled labels of the data.*

3.4 Discussion

We proposed a pMEM approach to quantify the dynamics of MEG oscillatory activity and applied this method to derive energy landscape measures, quantifying the abnor-

mal statistical characteristics of resting-state networks in JME patients. The number of within-network local minima from individual participant's energy landscape indicate the degree of multistability from an attractor network perspective [210] on MEG oscillatory power. The local minima are defined here, and should always be interpreted, in the context of a specific resting-state brain network. The energy values of minima on aggregated energy landscapes indicate the ease of transition from one stable state to another [389, 204], and its effects on network dynamics was demonstrated in the simulation of basin duration. Furthermore, the activation profiles of local minima provided key anatomical insights into functional configurations of cortical networks that differ between JME patients and controls. Our approach described network abnormalities in multivariate data from a statistical account. This extended previous research on the temporal evolution of system dynamics leading to seizures, which measures chaoticity [189] or entropy [369] in single or combined channels.

In this study, we found that patients with JME showed altered pMEM-derived energy landscapes in selective resting-state networks and frequency bands (Figure 3.7). For the energy landscapes estimated at the individual level, JME patients exhibited lower numbers of local minima than controls (Figure 3.4). For the aggregated energy landscapes estimated across participants, JME elevated relative energy values at the local minima of the FPN (theta, beta, and gamma bands) oscillatory activities (Figure 3.5). Our results confirmed the abnormalities of electrophysiological signals in JME [5], and provided new insights into JME pathophysiology affecting selective functional network configurations.

The fitted pMEM defined the energy values of all activity states of a network, from which an energy landscape of the network was depicted [389]. Because a local minimum of the energy refers to a network state with higher occurrence probability than its neighbouring states, the fewer number of local minima and elevated energy values in JME suggested alterations in the multi-stable dynamics of the brain networks that may be prone to perturbation and ictogenesis, in line with the dynamical disease account for epilepsy [111, 81, 373]. The energy landscape further allowed to characterise clusters of energy minima and their hierarchies in terms of the disconnectivity graphs (Figure 3.5). The disconnectivity graphs represent the ease of transition between local minima, i.e.

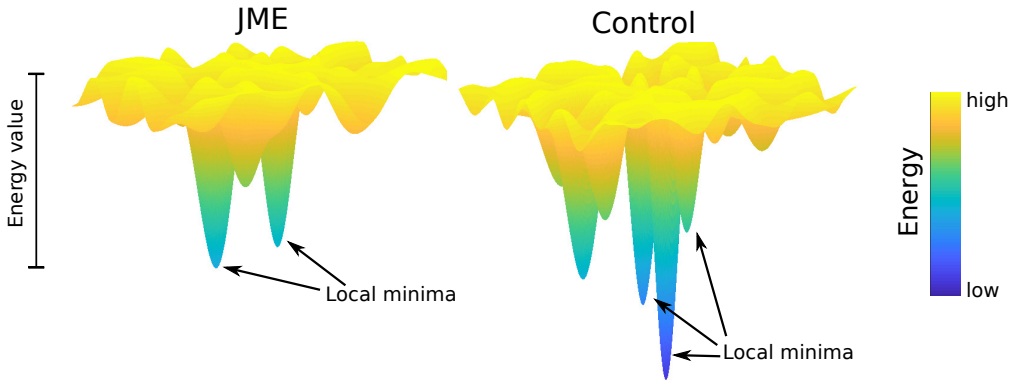

Figure 3.7: *A schematic diagram of altered energy landscape of MEG oscillatory power in JME patients (left) compared with controls (right). In selective functional networks and frequency bands, JME patients exhibited less local energy minima and elevated energy values (e.g. in theta-band FPN), suggesting that resting-state networks exhibit changes in the degree of multi-stability and in the ease of state transitions.*

energy barrier between the basins - the higher, the less probable is the switch between the states from these basins. In the FPN, the energy minima with bilateral high activation in the frontal or parietal regions were clustered separately and interleaved with lateralized energy minima (i.e. high activation in unilateral ROIs). This may indicate that network states with lateralized high activation represent transition states between frontal and parietal dominant states. In contrast, the DMN energy minima contained co-activation in bilateral ROIs, consistent with the evidence of strong interhemispheric and long-range connectivity in the DMN during awake [343, 20].

Our results highlighted JME as a distributed network disorder involving frontal and parietal lobes [285, 123, 452]. JME patients commonly exhibit impaired frontal cognitive functions [306], including working memory [388], decision making [466], response inhibition [218] and verbal fluency [294]. Demanding cognitive efforts during visuomotor coordination and decision-making can provoke myoclonic seizures in JME patients [458], and the degree of cognitive dysfunctions were associated with frontoparietal BOLD fMRI activity and connectivity [426]. Cortical and sub-cortical pathology may underlie the cognitive phenotype in JME. Activities in the lateral parietal cortex and precuneus have a dominant role in initiating and sustaining characteristic spike-

and-wave discharges in JME [235]. MR spectroscopy imaging of JME patients has identified reduced N-Acetyl aspartate concentrations in the frontal lobe and the thalamus [350, 473], which, together with widespread cortical morphological abnormalities [338], indicates dysfunctions in the corticothalamic loops in JME [167]. Further research should extend our results to associate specific abnormal energy minima to JME patients' cognitive and behavioural phenotypes.

We further demonstrated that the pMEM and energy landscapes can be used as a generative model to simulate the duration of the network activity in each energy basins (Figure 3.2) and as a predictive model for single-patient classification (Figure 3.6) beyond simple descriptive modelling [366]: it allowed us to combine measures from multiple energy minima to make inferences at an individual level. Such analysis, as demonstrated in the current study, would be useful in clinical applications for identifying patients from controls, or for detecting changes in electrophysiological data prior to seizure onset in future studies [369]. In addition, because classification-based analysis makes no assumption about data variances or distributions, it is a more stringent test than conventional statistical methods and provides accurate estimates of between-group differences [215]. The normalised energies of the theta-band FPN minima achieved the best classification results (>90%), comparable with other studies [145] and consistent with our hypothesis of selective abnormalities of oscillatory activity in JME. Indeed, pathological theta oscillation were reported as a hallmark of idiopathic generalised epilepsy [77], possibly owning to the involvement of the thalamus in initiating or facilitating theta oscillations through thalamocortical coherence [348].

The energy landscape measures for the SMN did not significantly differ between JME patients and controls. This result might seem counter-intuitive, given that motor cortex hyperexcitability has been reported in JME [18]. Nevertheless, previous research on resting-state functional connectivity also showed the lack of altered connectivity in the motor cortex in JME [244, 241, 115]. Our results suggested that the network states (i.e. patterns of co-activation) in the SMN, comprising pre- and post-central gyri as well as SMA, were not affected by JME during rest. However, this result does not rule out the possibility of network dysfunction in the motor circuit under stimulation or perturbation [426].

Our study provides new methods for studying the dynamics of MEG oscillatory activity. We showed that MEG oscillatory activity in resting-state networks was accurately described by the pMEM (Figure 3.3) and that the model fits were comparable between JME patients and controls. The pMEM was originally developed in the field of statistical mechanics and has been applied to population of spiking neurons [462]. More recently, it has been applied to quantify the dynamics of BOLD fMRI data [437, 436, 120, 13, 155]. However, achieving satisfactory pMEM fitting requires a large number of data samples [254, 389]. Because of the low temporal resolution of the BOLD signal, the applications of the pMEM to fMRI signals often need long scanning time that may be unrealistic for clinical populations, or to concatenate data across participants that limits the possibility of individual-level inferences [437, 13]. Here, we highlighted the feasibility and benefits of fitting the pMEM to MEG oscillatory power, which provided anatomically-specific and frequency-dependent results. Capitalizing on the high sampling rate of MEG, we showed that one can make inferences on energy landscapes at the individual level from a short recording session that was well tolerated by patients. Future studies could use longer recording sessions to systematically examine the effect of data length on pMEM fitting to MEG data.

Other methods are available to describe transient network dynamics. Microstate analysis from scalp EEG has identified successive short time periods during which the configuration of the scalp potential field remains semi-stable [20], and the spatial patterns of EEG microstates have been mapped onto distinct mental states [56, 267]. Recent studies using hidden-Markov models (HMM) characterized whole-brain spontaneous activity and identified hidden states with spatiotemporal patterns at durations of 100-200 ms [314]. Both microstate and HMM analyses are based on time-windowed approach and provide abstractions of the interactions within large-scale networks. In the current study, we defined the state of a network as an instantaneous snapshot of regional activities, and the pMEM provided a probabilistic model of the network states with minimum assumptions.

There are several limitations of this study. First, to quantify network dynamics as the occurrence probability of a finite number of network states, the oscillatory power in each ROI needed to be binarized (i.e. high vs. low activity), similar to other func-

tional connectivity studies [243]. The binarization procedure for applying pMEM in neuroscience differs between data modalities. For single unit recording and local field potentials (LFPs) [390, 462, 464], a threshold based on signal variance was applied to continuous data to identify active states (spikes in single units or negative deflections in LPFs). For resting-state fMRI data, a threshold based on the mean of BOLD responses was used [438, 204]. Unlike spiking trains or LFPs that have a clear definition of neuronal activity status, MEG oscillatory power reflects the level of synchronised activity in macroscopic neural populations, which, as a continuous measure, does not impose an a priori threshold for active/inactive binarization. The current study used the median of the oscillatory power envelope from each ROI as the threshold to binarize MEG source reconstructed data. The use of a median split is robust to signal outliers. Furthermore, our approach allows a common statistical criterion adaptive across regions and participants, appropriate for a potentially heterogeneous ensemble [88]. Future research could consider more complex quantification scheme such as ternary quantization that reduces oscillatory power to ternary values [475].

Second, the model fitting procedure for pMEM is computationally intensive. Currently, it is practically possible to fit the pMEM to a network of approximately 15 ROIs, because the number of network states increases exponentially with more ROIs. As a result, the current study focused on the dynamics within well-established large-scale resting-state networks, rather than a whole-brain network comprising all the regions. Other approximate model fitting procedures may allow us to extend our approach to larger networks with more ROIs [389], which is beyond the scope of the current study. Furthermore, the pMEM assumes the pairwise interaction between the brain states [389]. Although the previous studies showed successful applications of pairwise dynamics to the macroscopic networks [436, 120, 438], more and more evidence suggests that an interaction between brain areas might be asymmetric [187]. Yet, in order to reduce the complexity of the analysed system, it has been shown theoretically that reliable statistical models can be built without including all the interactions in a system; instead, pairwise interactions suffice [462, 389]. However, see also [340], where a criterion is derived for pairwise model generalisability to the realistic systems. Rather than focusing on dynamical changes of the epileptic system [247], this study analyses

resting state networks across JME patients and healthy controls. A question remains whether the bifurcation points or phase transition points in the pMEM method can predict the 'boundaries' between the two groups. Recently, an extension of the energy landscape framework was applied to study fluid intelligence. It was shown that human participants with higher IQ scores are closer to a critical state, i.e., the boundary between the paramagnetic phase and the spin-glass phase in Ising model [121]. Future research could inspect this question in the epilepsy setting. Moreover, to facilitate the new analyses, we have made our pMEM fitting procedure and analysis scripts open source and freely available (https://github.com/dokato/energy_landscape).

Third, the current study chose, *a-priori*, the AAL template for cortical parcellation. The AAL atlas is based on anatomical landmarks [413, 337] and commonly used in MEG resting-state analysis [302, 179, 341]. Previous studies have defined resting-state networks, including the ones used in our study, with the ROIs from the AAL atlas [339, 392]. It is worth noting that there is an abundant group of atlases for cortical parcellation with various levels of granularity [93, 95, 221, 151], and energy landscape measures from a network may change with different ROI definitions from an alternative atlas. Future research employing the pMEM for MEG need to make similar informed decisions on the choice of parcellation scheme based on specific research questions and intended networks. Also, the recent studies incorporate more data-driven splits of the brain's anatomy, e.g. by using hierarchical clustering [434], or cluster-permutation statistical methods [346].

Fourth, the sample size in the current study is sufficient for comparing and classifying between JME patient and control groups. However, JME often exhibits as a disease with a phenotypical spectrum, with variations among seizure frequencies, epileptic traits, and treatment response [28]. A larger clinical cohort with comprehensive neuropsychological assessments is necessary to investigate whether our energy landscape approach is sensitive to the quantitative spectrum of JME. For example, the impact of various drugs could be a confounding factor in this analysis. Hopefully, future studies with higher and more balanced sample of JME patients can help to address this issue. Moreover, the scanning protocol of this experiment did not enable continuous movement tracking. As a result, we could not directly compare the exact

level of movements between groups. Nevertheless, even if the residual movement artefacts in MEG data did differ between patients and controls, they would affect multiple networks and hence could not readily explain the network-specific group differences in energy landscape measures.

In conclusion, by fitting a pMEM to MEG oscillatory activity, we showed that JME patients exhibited atypical energy landscapes in selective brain networks and frequency bands, with a smaller number of local minima of the energy and elevated energy levels leading to altered multi-stable network dynamics. We further demonstrated that the pMEM and energy landscape offered generative and predictive power for discriminating between JME patients and controls. These results have the potential to be exploited in future diagnostic and pharmacological studies for a mechanistic understanding of ictogenesis in JME, for example by means of the large-scale simulations with neural-mass models [199], or DCM [220].

Chapter 4

Modelling neuronal populations: Integration of information in a perceptual decision-making task

In this Chapter two models are presented to characterise the neurocognitive process underlying perceptual decision-making with single or double information sources. We combined cognitive modelling and neural-mass modelling to explain the behavioural data from meticulously designed online experiment. Our modelling results showed that the addition of the second information source led to a lower signal-to-noise ratio of evidence accumulation with two congruent information sources, and a change in the decision strategy of speed-accuracy trade-off with two incongruent sources. Our findings support a robust behavioural change in relation to multiple information sources, which have congruency-dependent impacts on selective decision-making subcomponents.

4.1 Introduction

Making rapid decisions on the basis of noisy information is a hallmark of voluntary behaviour. Converging evidence from humans [170] and non-human primates [336] have supported an evidence accumulation framework governing the decision making process: the information is integrated over time until the accumulated information in

favour of one option reaches a response threshold. This integration process reduces the noise in the accumulated information and facilitates optimal behaviour in terms of accuracy and speed [47]. A large family of sequential sampling models [49] can describe adequately the cognitive [325] and neural [433] processes during decision making.

Much of research to date on simple decisions has focused on evidence accumulation from a single information source [148]. Understanding how a decision-maker integrates information from multiple sources is equally important. In preferential decisions with multiple attributes, such as buying a car based on its colour and price, sequential sampling models can effectively account for various biases and heuristics [62], supporting evidence accumulation as a parsimonious decision-making framework for distributed information sources.

The current study focused on another common scenario: making decisions by integrating the *same* type of information originated from multiple sources. For example, when approaching a T junction, a car driver has to consider incoming traffic from both left and right sides of the main road; while entering a roundabout, the driver only need to attend to one side because all vehicles circulate in one direction. An intriguing issue is: how does the presence of multiple information sources affect behavioural performance.

The progress of the understanding the decision making mechanism has been accelerated by the neuronal recordings from animal studies [273, 53, 148]. Typically they involve the paradigms from psychophysics, seeking the relation between sensory input and behaviour of the animal [159]. For example, a rhesus monkey is trained to indicate the dominant direction of movement of the dots visible on the screen - typically called random dot kinetogram (RDK). Typically - to simplify the problem - only two directions of the dots are considered. In a given trial, the coherence of dots can be controlled by the experimenter, who makes the task easier or mote difficult [282, 53]. The rhesus monkey responds to the stimuli with saccade eyes movement in the direction of choice in order to receive a reward. Selective cortical brain areas respond with increasing firing rates over time of recorded neurons before the decision is made. Current evidence suggests that the major areas engaged in the evidence accumulation are middle temporal visual area (MT) [282], superior colliculus (SC) [186], premotor cortex [175], caudate

nucleus [461] or lateral intra-parietal area (LIP) [217] (but see [207] where a behaviour of the monkey was intact after the inactivation of LIP), but also frontal areas of the brain: dorsolateral prefrontal cortex (dlPFC) [317, 170], frontal eye fields [41]. The animal studies were supported by the evidence from human research. For example, in [209] the amplitude of the centro-parietal positivity signal from the EEG recordings (in the medial parietal part of the scalp) was identified to correlate with the coherence strength of the dots movement. Several studies report involvement of a frontal-parietal network in the general process of perceptual decision making, which is complemented by the activity of deeper brain structures (in particular the basal ganglia) [211]. Also, fMRI research provided even more evidence that the activity within the left dlPFC is greater during easy decisions than during difficult decisions in the image categorization task [170, 171].

Apart from the source localisation of the neuronal decision correlates, several cognitive models emerged aiming to propose an algorithmic view on perceptual decision making. The majority of the models rely on the *sequential ratio analysis* procedure from the signal detection field. In the original setting, the log of ratio of probabilities of the observed evidence towards option 1 or 2 is collected with a commitment towards one of two hypothesis [431]. Thus, the procedure helps to convert a number of observations from the noisy sensory environment into a categorical choice [148]. In psychology, this test may be implemented in several ways, that have been shown to converge to just two: drift-diffusion models (DDM) [322, 325] and race models [414, 49]. Both introduce the notion of a decision variable that accumulates evidence towards one of the options. In the DDM this variable represents the difference between evidence towards each option [322]. In linear ballistic accumulator (LBA) model, which is a special case of a race model (other examples include a race diffusion model [399]), it is represented by two (for binary case) independent accumulators that race towards the threshold [58]. Typically, a decision model consists of the following parameters: non-decision time (representing sensory information encoding and movement), drift-rate (indicating a force with which a decision variable is attracted towards option), decision threshold (representing information accumulation stopping point), and bias (or preference towards one option) [49]. It has been shown that DDM shows an optimal behaviour

when its processing remains in a linear range [49]. The above models, even though, useful and conceptually simple do not tell much about the neuronal implementation of the underlying process.

Research on visual search provide circumstantial evidence to imply an imperfect integration of multiple sources, because of the limited capacity of the attentional system. When locating a target among similar distractors or filtering out task irrelevant information, there is a robust behavioural [299] and neural [331] cost in relation to selective attention on multiple sources. However, two important questions are yet to be addressed. First, when the total amount of information remains unchanged, does separating information into multiple congruent and incongruent sources have the same impact on behaviour? Second, does making decisions with additional information sources lead to a change in the speed of evidence accumulation, the decision threshold, or both?

Here, we addressed these questions in a pre-registered, carefully calibrated online experiment of perceptual decision in two independent groups. Human participants were instructed to decide the average motion direction of random-dot kinematogram from two tilted apertures (Figure 4.1). Coherent dot motion was presented in both apertures, with their moving directions to be congruent (both leftwards or rightwards) or incongruent (e.g. one leftwards and the other rightwards). In corresponding baseline conditions, coherent motion was presented in a single aperture, with the other aperture containing no coherent motion. Between the two groups, we varied the angular separation of the two apertures, allowing us to evaluate the repeatability of all within-subject effects and assess the between-group effect of aperture angles on behaviour. We fitted a cognitive model, DDM, to behavioural data and inferred the effects of motion congruency and sensory sources on model parameters. Furthermore, we extended a mean-field model (derived from biophysical model) [454] to demonstrate how the observed behavioural changes can be implemented by a biologically realistic neural network. Together, our study illustrated the neurocognitive mechanisms of perceptual decisions from multiple sources.

Table 4.1: *Statistical information about participants. NA - data not available.*

Category	Value
gender	female (25), male(67), NA(1)
handness	right(85), left(6), both(2)
age (years)	median: 25, mean: 27.3, std: 8.3
web browser type	Chrome(69), Internet Explorer(14), Firefox(4), Safari(5), NA(2)
nationality	United Kingdom(25), Poland(13), Portugal(11), United States(7), Spain(5), Italy(4), Mexico(3), Czech Republic(3), Denmark(2), Ireland(2), Hungary(2), France(2), Lithuania(1), Germany(1), Belgium(1), Sweden(1), Colombia(1), Estonia(1), Finland(1), Netherlands(1), Chile(1), Canada(1), Australia(1), South Africa(1), China(1), NA(1)

4.2 Methods

4.2.1 Participants and pre-registration

A total of 94 participants were recruited from an online recruitment portal (*Prolific*, prolific.co) and took part in the experiment online (age range 18-68 years old, median age 25 years old, 25 females, 85 right-handed). Table 4.1 summaries demographic features of the participants. All participants received monetary payments for their participation. Consent was obtained from all participants. We considered the recruitment from an online portal as a sample of convenience. The study was approved by the Cardiff University School of Psychology Research Ethics Committee.

Power analyses, exclusion criteria, experiment procedures and analysis plans were pre-registered prior to data collection (https://osf.io/4dn65). A sample size of $N > 44$ provides $> 90\%$ power to detect a medium within-group effect ($d = 0.5$) at $\alpha = 0.05$. We randomly assigned participants into two independent groups. Group 1 ($N = 49$) performed the perceptual decision task with two sources of visual inputs presented along $\theta = \pm 20°$, and Group 2 ($N = 45$) performed the task with visual inputs presented along $\theta = \pm 45°$ (see Procedure for details).

4.2.2 Apparatus

The experiment was conducted online. Experimental scripts for stimulus presentation and response collection were written in HTML with a JavaScript library jsPsych 6.0.5 [85] and the jspsych-rdk plugin [316]. The online experiment was hosted on a web server Pavlovia (pavlovia.org), and participants performed the experiment in web browsers on their computers (see Table 4.1). It has been shown that online experiments in modern web-browsers can serve as a suitable tool for measuring behavioural responses and reaction times with sufficient precision [86, 361, 11].

4.2.3 Stimuli

The visual stimuli contained two independent sets of random-dot kinematograms [54, 364, 264] displayed within two invisible rectangular apertures (140 pixels width, 550 pixels length) on a grey background (RGB=128, 128, 128). Both rectangular apertures located at the centre of the screen, with one tilted $+\theta$ from the horizontal plane and the other tilted $-\theta$. Hence, the two apertures formed an '$\times$' shape, with $\theta = 20°$ in Group 1 and $\theta = 45°$ in Group 2. To facilitate the integration of leftwards and rightward motion across apertures, four motion target indicators were presented at the end of the short edges of the two apertures. On each side of the screen (left or right), the two target indicators had the same colour (red or blue), and the colour assignment of those motion indicators was randomised across participants.

Each rectangular aperture contained 100 dots (i.e. 200 dots in total). Each dot had a radius of 3 pixels. We introduced coherent motion information along the long edge of each aperture (leftwards or rightwards). In each frame, a proportion of dots (namely the motion coherence) was replotted at an appropriate spatial displacement in the direction of motion (2 pixels/frame velocity), relative to their positions in the last frame, and the rest of the dots were replotted at random locations within the aperture. To minimise the impact of local motion information from individual dots, all dots were replotted at random locations after every 7 frames [316].

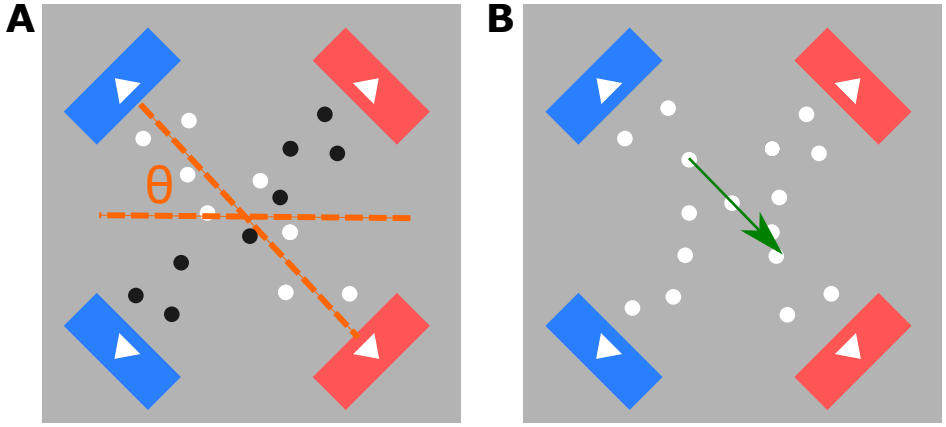

Figure 4.1: *The diagram of the RDK in two rectangular apertures used in (A) the staircase procedure and (B) the main experiment. θ denotes the angle between the aperture and the horizontal plane, which was $\pm 20°$ in Group 1 and $\pm 45°$ in Group 2. During the staircase procedure, one aperture contained block dots with 0% motion coherence. In the main experiment, both apertures contained white dots. Green arrow denotes the motion of the dots.*

4.2.4 Task and procedure

After informed consent and task instructions, the experiment included two parts: (1) a staircase procedure to identify two perceptual thresholds, and (2) the main perceptual decision-making task. In both parts, participants performed a two-alternative forced-choice (2AFC) task, deciding whether the coherent motion direction of the random-dot stimulus is leftward or rightward, either from a single information source (i.e. one aperture in Part 1, Figure 4.1A) or combined from double information sources (i.e. two apertures in Part 2, Figure 4.1B). Participants responded by pressing the 'k' key (for leftward decisions) or the 'p' key (for rightward decisions) on a keyboard with their right index and middle fingers. Participants took self-paced breaks after each part.

Part 1: staircase procedure

To allow participants to familiarise with stimuli and the task, participants underwent a short practise. The practise part consisted a single block of 32 trials. On each trial, one aperture contained black dots (RGB $= 255, 255, 255$) with 0% motion coherence, and the other aperture contained white dots (RGB $= 0, 0, 0$) at one of the four coherence

levels (5%, 10%, 20%, and 40%, 8 trials of each level). Participants were instructed to pay attention only to white dots (i.e. the informative aperture) and decide the direction of coherent motion. The coherent motion direction, the order of coherence levels and the informative aperture (i.e. the one at $+\theta^\circ$ or the one at $-\theta^\circ$) were randomised across trials. On each trial, the random-dot stimulus disappeared as soon as a response was made, or a maximum duration of 3500 ms was reached. The inter-trial interval was randomised between 900 and 1100 ms.

For online experiments, participants' hardware settings and their perceptual performance could vary substantially. Therefore, we measured motion discrimination thresholds using the same visual stimulus and the 2AFC task structure as in the practise: one aperture contained black moving dots with 0% (i.e. uninformative) coherence, and the other contained white dots with motion coherence set according to the staircase routine. The direction of coherent motion was randomised across trials. At the end of each trial, visual feedback in text was presented for 500 ms to indicate whether participant's response was correct or incorrect.

The staircase routine combined two parallel staircase procedures with fixed step sizes: one used a two-down/one-up rule and the other used a three-down/one-up rule. The two staircase procedures are independent and interleaved with each other. In both staircase procedures, the initial motion coherence was set to a supra-threshold value of 31.6%, the 'up' step size was 0.1 (log unit) and the 'down' step size was 0.074 [133]. Therefore, the two-down/one-up procedure converges to the coherence threshold of ∼71% accuracy (hereafter referred to as the low coherence c_{low}, and the three-down/one-up procedure converges to a coherence threshold of ∼83% accuracy (hereafter referred to as the high coherence c_{high}) [240, 134]. Each of the two staircase procedures terminated after ten staircase reversals and the corresponding threshold was calculated as the mean of the motion coherence levels for the last nine reversals.

Part 2: perceptual decisions from double sources

Part 2 is the main experiment, in which both apertures contained white dots (Figure 4.1B). Participants were instructed to attend to both apertures and decide whether

the coherent motion direction of all (white) dots was leftwards or rightwards.

After task instruction and a brief practise, the main experiment comprised 432 trials, which were divided into 6 blocks of 72 trials. Participants took self-paced breaks between blocks. Decision accuracy (proportion of correct responses) was measured after every two consecutive blocks. If a participant had the accuracy lower than 60%, the experiment ended prematurely and the dataset was discarded from further analysis.

Each block contained 50% of leftwards motion trials and 50% of rightwards motion trials. In each block, 64 main task trials and 8 control trials were presented in a randomised order. The task trials followed a 2 by 2 factorial design with two levels of motion coherence (high and low) and two levels of information sources (single source and double sources).

In the high coherence conditions, trials with double information sources had the low motion coherence c_{low} in one aperture and $c_{\text{high}} - c_{\text{low}}$ in the other. The coherent motion directions were *congruent* in the two apertures (i.e. both leftwards or both rightwards). Trials with single information source had the high motion coherence c_{high} in one aperture and 0% in the other.

In the low coherence conditions, trials with double information sources had the high motion coherence c_{high} in one aperture and $c_{\text{high}} - c_{\text{low}}$ in the other. Importantly, the coherent motion directions were *incongruent* (i.e. opposite) in the two apertures. Trials with single information source had the low motion coherence c_{low} in one aperture and 0% in the other. Therefore, for both double and single information sources, the net motion coherence was always c_{high} in high coherence conditions and c_{low} in low coherence conditions. In control trials, the motion coherence levels in two apertures was set to 60% and 0%. These easy control trials were served as attention check and excluded from subsequent data analyses.

Each trial started with a 250 ms fixation period, during which a black cross presented in the central of the screen. RDK stimuli in two apertures were then presented for a maximum period of 4000 ms, and the stimuli disappeared as soon as a choice was made. The visual stimulus was followed by an inter-trial interval randomised between 400 and 600 ms.

The example of the stimuli from both parts of the experiment is demonstrated in

the following video: https://youtu.be/mhvpiYLXHag.

4.2.5 Data analysis

For the staircase procedure, the non-parametric Mann-Whitney U-test was used to compare the high and low coherence levels (c_{high} and c_{low}) between the two aperture angles ($\theta = 45°$ and $\theta = 20°$). 95% confidence intervals (CI) were obtained using bootstrap procedure with 1000 resamples of simulated distributions.

For the main experiment, we quantified response time (RT) of each trial as the latency between the RDK stimulus onset and behavioural response. To eliminate fast guesses, trials with RT faster than 250 ms were removed. Trials without a valid response were also removed. The discarded trials accounted for 0.26% of all trials. We used mixed frequentist and Bayesian ANOVAs to make group inferences on mean decision accuracy and RT, with the coherence level and the number of information source as within-subject factors. Assumptions of variance equality were checked with Levene's test. We performed post-hoc comparisons using JASP (jasp-stats.org) and used Bayes Factors (BF_{incl}, BF_{10}) to characterise the strength of evidence [429].

4.2.6 Cognitive modelling of behavioural data

We used the hierarchical DDM toolbox [447] to fit DDMs to individual participant's response time distribution and decision accuracy. The hierarchical DDM assumes that the model parameters of individual participants are sampled from group-level distributions, and the Bayesian fitting procedure estimates the posterior distributions of all model parameters at both individual and group levels, given the observed data.

The basic form of the DDM contained three core parameters [325]: (1) the drift-rate (v), (2) the decision threshold (a), and (3) the non-decision time (T_{er}). For each trial, the model assumes that noisy information is accumulated over time at an averaged rate of v and a starting point of $a/2$, until the accumulated information reaches the upper or the lower decision boundary (a or 0) that indicates a correct or incorrect binary response, respectively. The model prediction of RT is the sum of the duration of the accumulation process and the non-decision time, with the latter accounting for delays

in sensory encoding and motor execution [206].

To accommodate changes in behavioural performance between conditions, we estimated four variants of the DDM with different parameter constraints. The first three variants allow two of the three parameters (v, a, T_{er}) to vary between conditions, and the last variant allows all three parameters to vary. All parameters are allowed to vary between participants in all variants.

For each variant, we generated 20,000 samples from the joint posterior distribution of all model parameters by using Markov chain Monte Carlo sampling. The initial 4000 samples were discarded for the sake of obtaining stable posterior estimates [447]. To improve the model's robustness to outliers, we estimated mixture models, in that 95% of the data are explained by the DDM, and 5% of the data are expected to be outliers generated from a uniform distribution [328].

Model fits were assessed by comparing each model's deviance information criterion (DIC) value [371], which takes into account both the log-likelihood function of observed data and the complexity of the model. For the best fitting variant, we used Bayesian hypothesis testing [139] to make inferences between conditions from the parameters' group-level posterior distributions. For consistency, we use p to refer to frequentist p-values, and $P_{p|D}$ to refer to the proportion of posteriors supporting the testing hypothesis at the group level from Bayesian hypothesis testing.

4.2.7 Recurrent neural mass model

We further used a mean-field model [454] to qualitatively demonstrate the effects of motion coherence and the number of information source on behavioural performance. The model considered here is simplified from a recurrent spiking neural network model [433] via the mean-field approximation (as described in Section 1.4.2). Specifically, the neural-mass model includes two neural populations (i.e. accumulators) supporting two alternatives in a decision-making task. The neural accumulators have self-excitatory and mutual inhibitory connections. Each accumulator receives selective external inputs ($I_{\mathrm{in,L}}$ and $I_{\mathrm{in,R}}$) as momentary evidence supporting each alternative (e.g. leftwards vs rightwards motion), as well as a common, non-selective background input I_0 (Fig-

ure 4.7A). During decision making, two accumulators compete against each other, and the first accumulator that reaches a decision threshold renders the corresponding response. It has been shown that this biologically realistic neural mass model can explain behavioural and single-unit recording data from perceptual decision experiments [454]. Moreover, within a certain parameter range, the dynamics of the model can mathematically approximate that of the DDM [454, 49].

Here, we extended the original neural-mass model to take into account the presence of the two information sources in the current study (for modelling details see Supplementary methods in Appendix B). The deterministic input currents ($I_{\text{in,L}}$ and $I_{\text{in,R}}$) to the two neural accumulators are given by

$$\begin{cases} I_{\text{in,L}} = J_{\text{ext}}[\alpha\mu(1+c_1) + (1-\alpha)\mu(1+c_2)] + \beta I_0, \\ I_{\text{in,R}} = J_{\text{ext}}[\alpha\mu(1-c_1) + (1-\alpha)\mu(1-c_2)] + \beta I_0, \end{cases} \tag{4.1}$$

where the first term is the selective input current and the second represents the non-selective background input current. Here, c_1 and c_2 denote the motion coherence levels in the two independent apertures. For simplicity, hereafter we assign c_1 to represent the stronger coherence between the two ($|c_1| > |c_2|$). Other parameters were set in line with previous studies [454, 375]: $J_{\text{ext}} = 5.2 \cdot 10^{-4}$ nA·Hz^{-1} is the the average synaptic coupling parameter, $I_0 = 0.321$ nA represents the baseline of the background input current, and $\mu = 35$ Hz is the baseline of input strength of the evidence.

The two scaling parameters α and β control to what extent task conditions affect model inputs. First, in trials with double informative sources, participants need to combine the evidence from two apertures for optimal decisions. The parameter α determines how a decision maker splits the weight of sensory evidence from two sources. $\alpha = 0.5$ implies that a participant weight two sources equally, while $0.5 < \alpha < 1$ or $0 < \alpha < 0.5$ implies that the dominate source is weighted more or less, respectively.

Second, previous studies suggest that a change in the baseline input I_0 results in speed-accuracy trade-off [375, 173]. Compared with the single source condition, the double source condition with incongruent motion directions had lower accuracy and faster RT, suggesting that participants may trade accuracy for speed in the presence of

conflict information. Therefore, we assumed that the non-selective background input current is modulated by a factor of β in that condition, which changes the model dynamics and in turn affects both the accuracy and RT relative to the condition with a single information source. For other conditions, we set $\beta = 1$ such that the non-selective input is at its baseline level.

To identify the parameter regime where the neural mass model can produce qualitatively the behavioural pattern observed in the experiment, we ran model simulations with different values of α and β. For each parameter set, we ran 8,000 simulations of the four experimental conditions with representative coherence levels ($c_{\text{high}} = 20\%$ and $c_{\text{low}} = 15\%$). For example, in the incongruent condition with double information sources, $c_1 = 20\%$ and $c_2 = -5\%$ for leftward motion; and $c_1 = -20\%$ and $c_2 = 5\%$ for rightward motion. Mean accuracy and RT of each condition was then calculated from all simulations.

4.2.8 Open data and scripts

We have made the data (https://figshare.com/articles/dataset/13567916), all analyses scripts and experimental materials (https://github.com/dokato/2drdk) open access.

4.3 Results

4.3.1 Behavioural results

Two groups of participants performed a coherent motion discrimination task, with independent RDK stimuli in two apertures at an angle of 20° (Group 1) or 45° (Group 2) (Figure 4.1A). Prior to the main experiment, all participants underwent a fixed-size staircase procedure to estimate two motion coherence thresholds: c_{low} from a two-down/one-up rule and c_{high} from a three-down/one-up rule (Figure 4.2 and Supplementary Figure B.1). A Mann Whitney U test showed no significant difference in coherence thresholds between the two groups with different aperture angles (c_{low}: $U(45, 49) = 983$, $p = 0.18$, 95%CI $= [-2.7, 0.9]\%$; c_{high}: $U(45, 49) = 952$, $p = 0.13$,

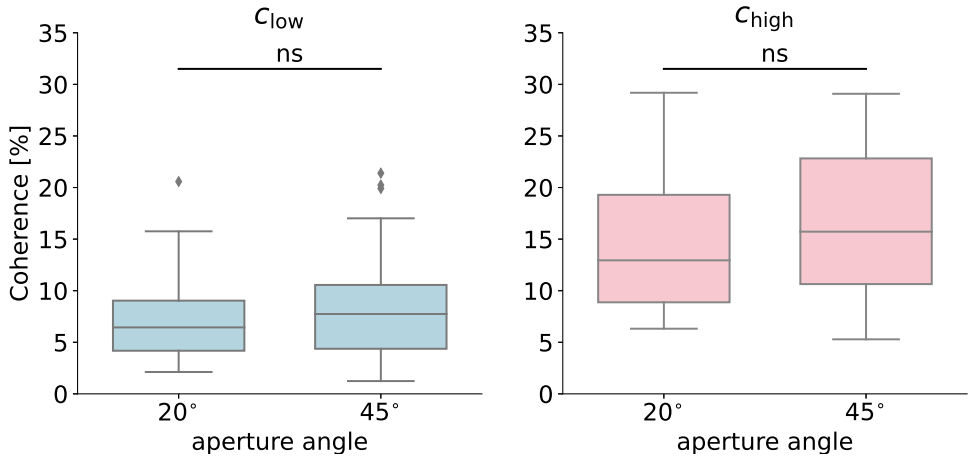

Figure 4.2: *Staircase procedure results. Two motion coherence thresholds obtained from a parallel staircase routine: c_{low} from the two-down/one-up rule and c_{high} from the three-down/one-up rule. There were no significant (ns) differences between coherence values for the aperture angles $\theta = 20°$ and $\theta = 45°$.*

95%CI $= [-5.4, 1.3]$%). As expected, the coherence threshold c_{high} was significantly larger than c_{low} ($U(94) = 1534$, $p < 0.0001$, 95%CI $= [-9.1, -5.1]$%). These results suggest that participants achieved reliable performance in the motion discrimination task online.

In the main experiment, participants in both groups decided the combined coherent motion direction (leftwards vs. rightwards) in a 2-by-2 factorial design: either single or double apertures contained non-zero motion coherence, and the combined coherence level in the two apertures was either c_{low} or c_{high} (Figure 4.1B). We quantified participant's performance in mean decision accuracy (proportion of correct) and RT.

The two groups with different angles of stimulus apertures achieved similar performance. A two-way mixed ANOVA showed no significant group effect on behavioural performance (accuracy: $F(1, 92) = 0.009$, $p = 0.92$, $\eta_p^2 < 0.001$, $\text{BF}_{\text{incl}} = 0.29$; RT: $F(1, 92) = 0.05$, $p = 0.30$, $\eta_p^2 < 0.001$, $\text{BF}_{\text{incl}} = 0.24$). The participant's grouping interacted with combined coherence levels for accuracy: ($F(1, 92) = 4.27$, $p = 0.04$, $\eta_p^2 = 0.04$, $\text{BF}_{\text{incl}} = 4.20$), but not RT: $F(1, 92) = 1.04$, $p = 0.31$, $\eta_p^2 = 0.01$, $\text{BF}_{\text{incl}} = 0.19$), nor the number of information sources (accuracy: $F(1, 92) = 0.02$, $p = 0.88$, $\eta_p^2 < 0.001$, $\text{BF}_{\text{incl}} = 0.17$; RT: $F(1, 92) = 0.01$, $p = 0.93$, $\eta_p^2 < 0.001$,

$\mathrm{BF}_{\mathrm{incl}} = 0.15$).

Across both groups, the high combined coherence c_{high} led to better accuracy (Figure 4.3, $F(1, 92) = 269.25$, $p < 0.001$, $\eta_p^2 = 0.745$, $\mathrm{BF}_{\mathrm{incl}} = 2.3 \cdot 10^{55}$) and faster RT (Figure 4.4, $F(1, 92) = 53.70$, $p < 0.001$, $\eta_p^2 = 0.37$, $\mathrm{BF}_{\mathrm{incl}} = 3.73$) than the low combined coherence c_{low}. Compared with conditions of single information source, splitting motion information into two apertures resulted in lower accuracy ($F(1, 92) = 47.50$, $p < 0.001$, $\eta_p^2 = 0.34$, $\mathrm{BF}_{\mathrm{incl}} = 1.7 \cdot 10^4$) without a significant main effect on RT ($F(1, 92) = 0.14$, $p = 0.71$, $\eta_p^2 = 0.002$, $\mathrm{BF}_{\mathrm{incl}} = 0.12$).

There was no interaction in accuracy between combined coherence levels and the number of information sources ($F(1, 92) = 0.01$, $p = 0.92$, $\eta_p^2 < 0.001$, $\mathrm{BF}_{\mathrm{incl}} = 0.17$). For RT, the interaction between the two factors was significant ($F(1, 92) = 208.38$, $p < 0.001$, $\eta_p^2 = 0.69$, $\mathrm{BF}_{\mathrm{incl}} = 1.2 \cdot 10^{13}$), suggesting that presenting motion information in double apertures elicited different changes in response speed between the two combined coherence levels (Figure 4.4). It is worth noting that in c_{high} conditions, motion directions in two apertures were congruent in the case of double information sources (i.e. both leftwards or both rightwards). A post-hoc test showed that compared with single source conditions, congruent double source conditions had slower RT ($\mathrm{BF}_{10} = 8014.01$, Bayesian t-test). Conversely, in c_{low} conditions, motion directions in two apertures were incongruent (e.g. one leftwards and the other rightwards) and led to faster RT ($\mathrm{BF}_{10} = 1.21 \cdot 10^4$).

4.3.2 Cognitive modelling results

We used a hierarchical Bayesian implementation [447, 419] of the DDM [322, 49] to decompose individual participant's accuracy and RT into model parameters that quantify latent cognitive processes. We considered four model variants, which allow the drift-rate v, the non-decision time T_{er} and the decision threshold a to be fixed or vary between task conditions.

For each model variant, the Gelman-Rubin $\hat{R}$ convergence criterion [137] was used to assess the convergence of the last 16,000 MCMC samples from 5 independent Markov chains. The maximum value of the statistic from all parameters was $\hat{R} = 1.0012$,

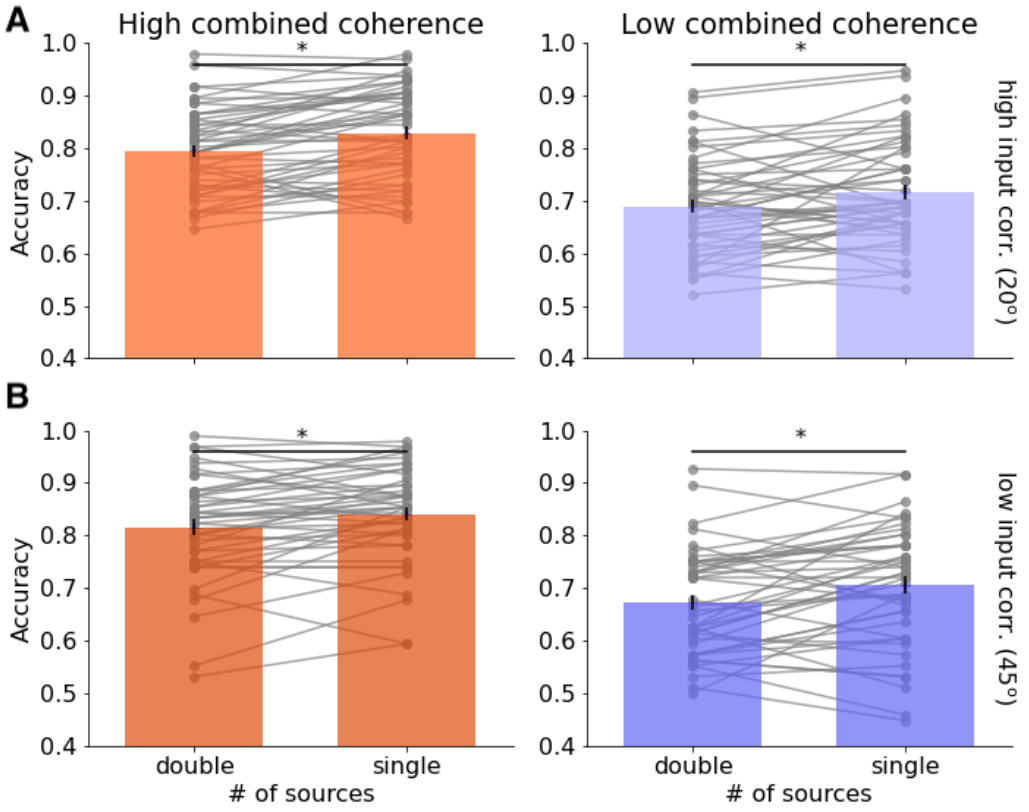

Figure 4.3: *Accuracy (proportion of correct responses) in the main experiment for high (red) and low (purple) combined evidence condition. Bars represent the averaged accuracy in (A) Group 1 (aperture angle $\theta = \pm 20°$) and (B) Group 2 ($\theta = \pm 45°$). Grey dots represent individual participants' accuracy. Each solid line links the performance between double and single source conditions from the same participant. Solid horizontal line with asterisk indicate the significant pairwise comparison. Solid vertical lines represent the standard error.*

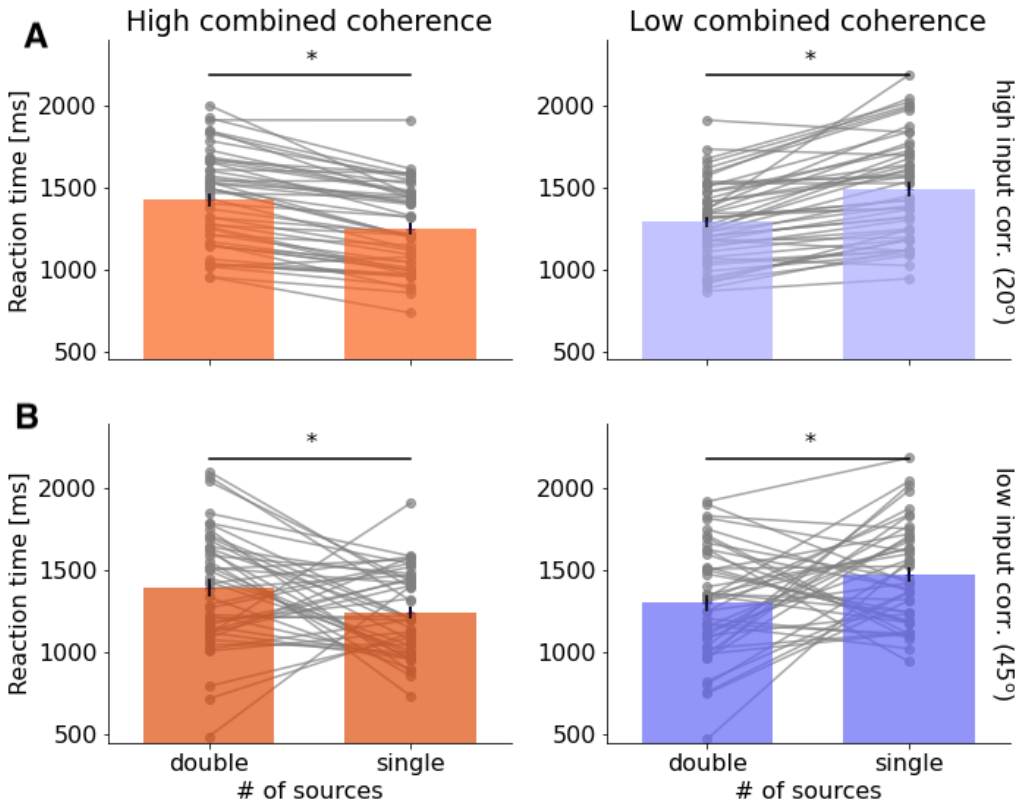

Figure 4.4: *Reaction time (RT) in the main experiment for high (red) and low (purple) combined evidence condition. Bars represent the averaged accuracy in (A) Group 1 (aperture angle $\theta = \pm 20°$) and (B) Group 2 ($\theta = \pm 45°$). Grey dots represent individual participants' accuracy. Each solid line links the performance between double and single source conditions from the same participant. Solid horizontal line with asterisk indicate the significant pairwise comparison. Solid vertical lines represent the standard error.*

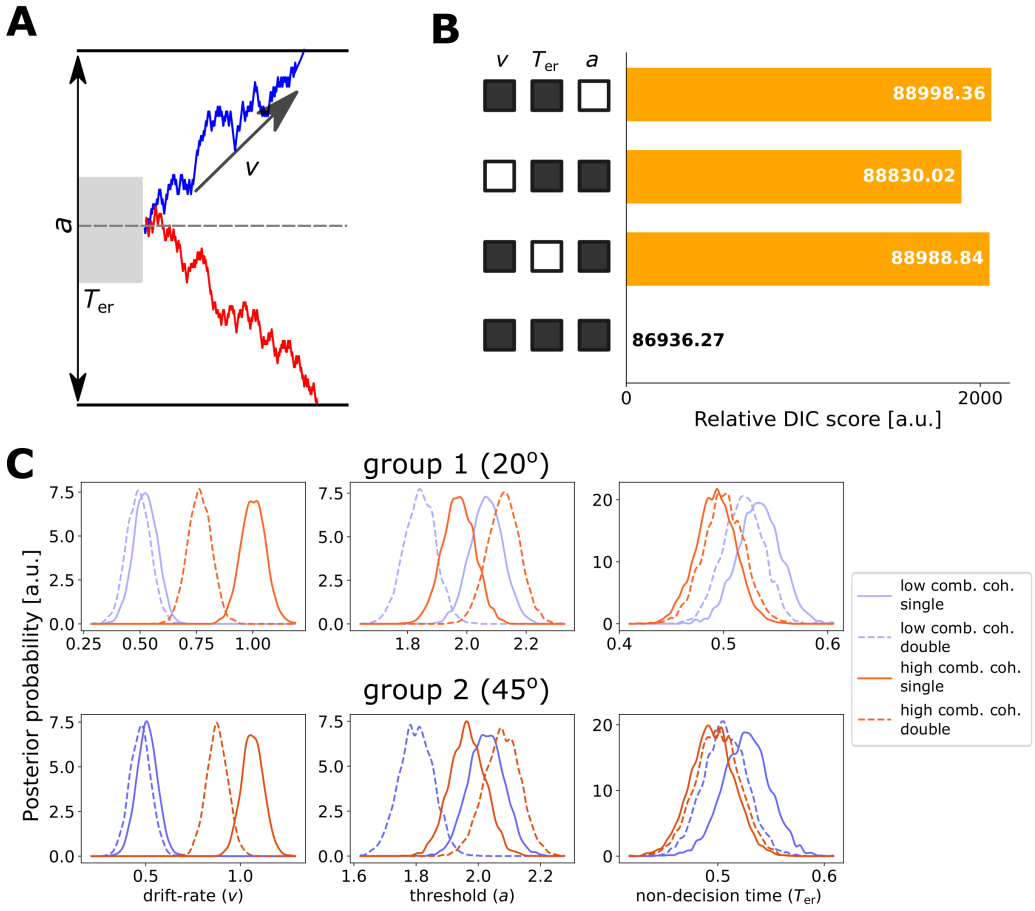

Figure 4.5: *Drift-Diffusion Model (DDM) fitting results. (A) Examples of evidence accumulation trajectories depicted by the DDM. The decision threshold a represents the distance between the correct and incorrect decision thresholds. The drift-rate v describes the average speed of evidence accumulation. The non-decision time* T_{er} *represents the latency of other processes not included in the evidence accumulation. The diffusion continues until the accumulated evidence reaches one of the two thresholds (solid black lines). If the accumulated evidence reaches the correct (upper) threshold (blue trajectory), the model predicts a correct response. Because of noise, the accumulated evidence may reach the incorrect (lower) threshold (red trajectory). (B) The deviance information criterion (DIC) value differences between the four variants of the DDM and the best fit. The black square indicates that the corresponding parameter can vary between the conditions, and the white square indicates that the parameter is invariant. The best model had variable a, v and* T_{er} *between conditions. (C) The posterior distributions of parameter values of the best fit model (top: Group 1 with* 20° *aperture angle; bottom: Group 2 with* 45° *aperture angle).*

Table 4.2: *Posterior comparisons of model parameters. The table listed the proportion of non-overlap between two posterior parameter estimates x and y, which is equivalent to a Bayesian test of the hypothesis $P_{p|D}(x > y)$. Experimental conditions: I1 and I2 refer to single or double informative sources with a combined coherence of c_{high}; C1 and C2 refer to single or double informative sources with a combined coherence of c_{low}. The DDM model parameters: v drift-rate, a decision threshold and T_{er} non-decision time. Two angular distances in two groups: $\theta_1 = 20°$, $\theta_2 = 45°$.*

	v		a		T_{er}	
	θ_1	θ_2	θ_1	θ_2	θ_1	θ_2
x = I1; y = I2	0.662	0.671	0.997	0.998	0.708	0.744
x = C1; y = I2	1.000	1.000	0.958	0.984	0.173	0.339
x = C2; y = I2	1.000	1.000	0.999	0.999	0.250	0.402
x = C1; y = I1	1.000	1.000	0.113	0.233	0.067	0.145
x = C2; y = I1	0.999	1.000	0.770	0.734	0.118	0.186
x = C1; y = C2	0.998	0.989	0.025	0.087	0.392	0.430

which is lower than the criterion of convergence 1.1 [137], suggesting that all parameter estimates converged after 20,000 steps.

The model variant that described the data best (i.e. the one with the lowest DIC value) allows all three parameters (v, T_{er} and a) to vary between conditions. To evaluate the model fit, we generated model predictions by simulations with the posterior estimates of the model parameters. There was a good agreement between the observed data and the model simulations in all conditions (Figure 4.6 and Supplementary Figures B.2 and B.3).

Figure 4.5C shows the posterior parameter estimates for the two participant groups. We used Bayesian statistics [139, 231] to quantify the proportion of parameters' posterior distributions that did not overlap between groups and conditions (Table 4.2). There was no evidence to support a difference in model parameters between groups ($P_{p|D} < 0.93$ in all parameters). For the drift-rate, there were strong evidence to support differences between all conditions (Group 1: $P_{p|D} > 0.998$; Group 2: $P_{p|D} > 0.989$, Table 4.2), except between the single and double source (i.e. the incongruent condition) conditions with the low coherence level (Group 1: $P_{p|D} = 0.662$; Group 2: $P_{p|D} = 0.671$). The incongruent condition with double information sources had a lower decision threshold than the other three conditions (Group 1: $P_{p|D} > 0.958$; Group 2:

$P_{p|D} > 0.984$). We did not observe strong evidence in supporting a difference in the non-decision time between conditions (Group 1: $P_{p|D} < 0.71$; Group 2: $P_{p|D} < 0.74$).

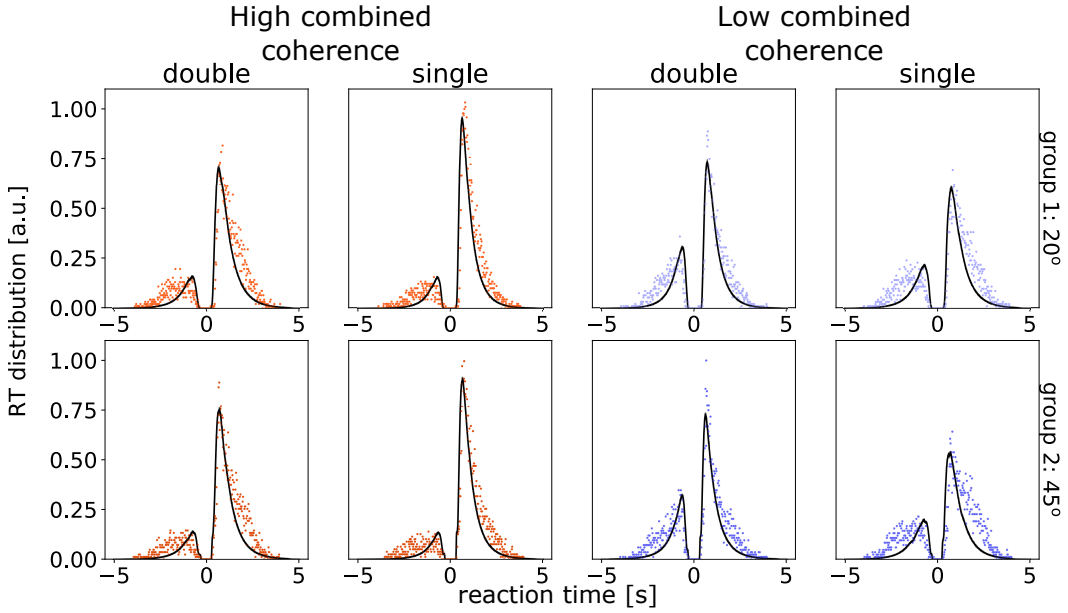

Figure 4.6: *Posterior predictive response time (RT) distributions from the fitted DDM. Each panel shows normalised histogram of the observed data (red for high combined coherence and blue for low combined coherence conditions) and the model predictions (black lines) across participants. The RT distribution of correct responses is shown along the positive horizontal axis. The RT distribution of error responses is shown along the negative horizontal axis. The posterior predictions of the model were generated by averaging 5000 simulations of the same amount of observed data. Top row shows results from Group 1 (with* 20° *angle between apertures) and the bottom row for Group 2 (*45°*).*

4.3.3 Neural-mass modelling results

Our cognitive modelling results suggested that splitting coherent motion information into two apertures led to a decrease of drift-rate in the congruent condition, and a decrease of decision threshold in the incongruent condition. How could these changes be incorporated in a biologically realistic model? To address this question, we introduced two extensions (Figure 4.7A) to a neural-mass model of perceptual decision [454], which implements an evidence accumulation process akin to that of the DDM [49]. First, for conditions with double information sources, we assumed that the sensory input selective

to motion coherence is a weighted sum of the two sources. The two weights ($\alpha\mu$ and $(1-\alpha)\mu$; see Equation 4.1) sum up to the constant baseline weight μ that is applied to the conditions with single information source. Second, we assumed that the non-selective sensory input I_0 is changed at the rate of β in the double source condition with incongruent inputs, which has been shown to be a realistic neural mechanism in modulating decision threshold [173, 374].

We parametrically modulated the two scaling parameters α and β. For each parameter set, we simulated the extended neural-mass model 20,000 trials (5,000 simulations for each experimental condition) and estimated the decision accuracy as well as mean RT. Figures 4.7B and 4.7C show the behavioural performance from simulations. We further identified parameter regimes that qualitatively satisfy the observed performance difference between double (both congruent and incongruent) and single information sources conditions. Based on model simulations, the scaling parameter α on input weights needs to be larger than 0.5, suggesting that the dominate, or more informative, sensory input of the two apertures is weighted more than the other. The parameter β needs to be larger than 1, suggesting that incongruent double information sources are associated with an elevated non-selective sensory input.

4.4 Discussion

Integration of information across multiple channels by human decision makers has been shown to be imperfect in many experimental settings [365]. Traditional theories link such limitations to an information bottleneck, which gates entry into central decision systems (eg. basal ganglia) [55]. Early evidence from psychology of sensation posited that multiple information sources affect the judgement [196]. This was hypothesised to be related to the attention mechanism, affecting the *coding stage* of the decision making process [365]. More recent studies, confirmed the neuronal correlations of such explanation [331]. In [331], the authors showed that qualitatively different firing rates (from single-unit recordings of macaque's brain) depend on whether a stimulus appeared alone or accompanied by distractors.

The current study examined, in two independent groups, how the presence of a sec-

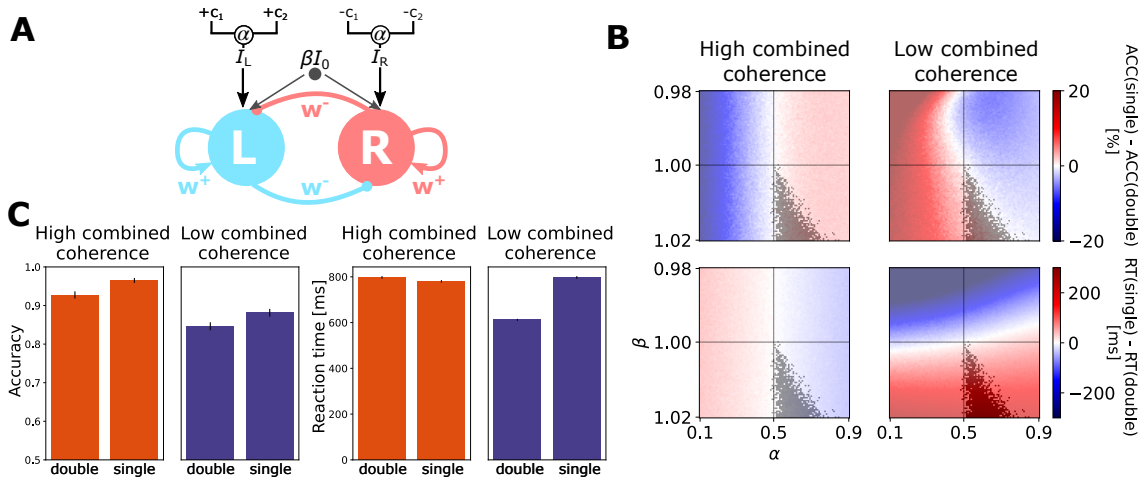

Figure 4.7: *Neural-mass model simulation results. (A) The diagram of the two-state neural-mass model. w^+ denotes excitatory connections, w^- denotes inhibitory connections, α determines how the weight of sensory evidence is split between two sources, and β modulates non-selective background input current. (B) Parameters space with difference in performance between single and double informatino sources in high combined evidence (left) and low combined evidence (right) conditions in terms of accuracy (top row) and reaction times (bottom row). Grey zone indicates area where the model parameters reproduce the direction of behavioural differences observed in the experiment. α varied between 0.1 and 0.9, and β varied between 0.98 and 1.02. (C) Behavioural performance from model simulations with parameters $\alpha = 0.7$ and $\beta = 1.018$.*

ond source of sensory information affects the behavioural performance of perceptual decision as well as its underlying neurocognitive mechanisms. When motion directions are congruent between the two sources, decisions on the global motion direction were less accurate and slower than that in the single-source condition with the same amount of total information (i.e. combined motion coherence). In contrast, when two information sources are incongruent, decisions were less accurate but faster than that in the single-source condition. Therefore, the change in task performance depends on the congruency between multiple sources of sensory evidence.

Using a Bayesian DDM, our cognitive modelling provided novel evidence on the decision-making process with multiple information sources. First, information congruency has selective influence on different decision making subcomponents. The congruent, double-source condition had a lower drift-rate than its corresponding single-source condition (i.e. with a combined motion coherence of c_{high} in both). The drift-rate of the DDM represents the signal-to-noise ratio of the information [325] and has been linked to the allocation of attention [354]. The presence of congruent information in two apertures may modulate the divided attention towards the stimuli that in turn lowers the averaged rate of evidence accumulation. Human electrophysiological data support this proposition. For two peripheral visual patches presented simultaneously, the early-visual ERP component is characteristic to the attended location, as its amplitude is maximal over posterior electrodes contralateral to the attending side [108, 251]. Recent studies showed that this EEG marker of selective attention modulates the rate of evidence accumulation in perceptual decision [249], and the dynamics of selective attention can influence evidence accumulation throughout the decision process [319].

Second, splitting motion information into two incongruent apertures did not vary the drift-rate. Instead, there was a substantial reduction in the decision threshold, reflecting the behavioural change that participants traded accuracy for speed in this condition. The speed-accuracy trade-off (SAT) is widely observed across decision making tasks [446, 172, 31]. In experiments with humans, the SAT is often induced explicitly via verbal instructions [469] or response deadlines [459]. Such manipulations can efficiently switch between accuracy-seeking and speed-seeking behaviour every few trials [275] or in consecutive trials [128]. Modelling studies on explicit SAT demands have

been consistently associated with the change of decision threshold [300, 323]: a smaller decision threshold leads to faster and more error-prone decisions. Nevertheless, the SAT can also be triggered endogenously without explicit demands [92]. The incongruency of the information might affect a certainty of choice. Recent mean-field modelling studies suggested that the level of certainty modulates the response times [15]. In the proposed model, an uncertainty-monitoring module receives excitatory input from the sensorimotor module, similar to the model proposed in this study, and provides to in an excitatory feedback with an intermediate inhibitory connections. The uncertainty-encoding population receives a constant tonic excitatory input that varies across trials in specific cases. This module was based on transient neural dynamics observed in animal and human studies [15]. In the current study, the two apertures in the incongruent condition contained contradictory information, presenting a decision dilemma. Our results showed that, in such difficult scenario, participants adapted their decision strategy to be more speed-seeking, allowing them to complete the current decision sooner. Future research could examine this conflict avoidance bias further, by changing the relative difference between multiple incongruent information sources.

Third, it is worth comparing between single- and double-source conditions which had equal motion coherence in the dominate aperture. Compared with the single-source condition with high coherence (c_{high} in one aperture and 0% in the other), the incongruent double-source condition (c_{high} in one aperture and $c_{\text{high}} - c_{\text{low}}$ in the other) condition had a smaller drift-rate. The congruent double-source condition had a larger drift-rate than the single-source condition with low coherence. That is, introducing additional incongruent (or congruent) information led to a reduction (or increase) in the rate of evidence accumulation. These results agree with two robust behavioural effects consistently reported in the literature of visual search: the presence of distractors in hindering the search performance [299], as well as the facilitating role of task relevant information [230]. Our findings further support that participants did not only attend the dominate aperture, but attempt to integrate motion information across apertures to form decisions, albeit the integration of multiple information sources was not optimal, as discussed above and reported elsewhere [457].

Fourth, the T_{er} is considered as the latency external to the evidence accumulation

process [325]. Recent electrophysiological and imaging studies suggest that the T_{er} accounts for delays in early sensory processing [289] or motor preparation [206]. The current study did not observe a change in the T_{er} between task conditions in either participant group. Hence, our results are unlikely originated from potential changes in early visual processing or motor execution in response to multiple information sources.

Based on our cognitive modelling, we proposed two extensions to a neural-mass model of decision making [454]. The first extension is to vary the baseline weights between sensory inputs from two independent apertures, and the second is to vary the non-selective background inputs in the incongruent double source condition. From an exhaustive search of the parameter space, we identified the parameter regime that can qualitatively account for the observed behavioural changes in the presence of two information sources. It is worth noting that the neural-mass model is not meant to fit to experimental data, but provides a biologically plausible interpretation of their neural implementations.

We showed that, to accommodate experimental results, the sensory input from the dominant source needs to be weighted higher than the input from the additional source ($\alpha > 0.5$). When this ratio becomes too high, the contribution of the additional source diminishes, resulting in the model unable to integrate information from the non-dominant source. Therefore, perceptual decisions with two information sources involve an unbalanced integration that is biased towards the more informative source. Although this study does not attempt to cover the neuronal mechanism of attentional effects, the previous modelling work showed that neuronal correlates of attention can be achieved by simultaneously increasing the gains of both excitatory and inhibitory neurons [286]. This agrees with the experimental evidence that gain modulation of neuronal firing rates - dependent on behavioural context and attention - is found in the parietal and extrastriate cortical areas [406].

Additionally, the non-selective background input needs to be elevated in the incongruent double-source condition ($\beta > 1$). An increased baseline activity effectively decreases the amount of evidence required to reach a decision threshold [374], leading to speed-seeking behaviour at the cost of less accurate decisions that was observed in the current study. Both brain imaging [191] and single-unit recording [173] studies

showed that the baseline change underlies the SAT, consistent with our model simulation results.

Interestingly, although participants were instructed to decide leftwards vs. rightwards coherent motion from two tilted apertures, the angular distance between the apertures did not affect behaviour nor DDM parameters. This may seem counterintuitive, because a larger angular distance results in less coherent motion information to be projected onto the horizontal plane. Future studies could examine whether there is a significant behavioural difference at larger aperture angles, because, in an extreme condition of two vertical apertures ($\theta = \pm 90°$), there is zero horizontal motion and the decision accuracy will be at chance. One plausible account for the lack of group difference is that participants decided the coherent motion direction with a reference of individual apertures (i.e. along their long edges), not the horizontal plane. One could validate this hypothesis by presenting multiple independent sources of motion information within a single aperture (e.g. [439]).

There are several limitations of this study. First, as in all online experiments, the current study faced practical constraints that could affect the millisecond-level precision of stimulus timing [11]. To mitigate the impact of variable testing environments between participants, we pre-registered the experiment, applied rigorous inclusion/exclusion criteria, conducted staircase procedures to calibrate stimuli for individual participants, and focused on within-subject effects in most analyses. Our study and research practises contribute to the growing trend of online psychological, or even psychophysical experiments, confirming the feasibility and reproducibility (i.e. in two independent groups) of online experiments to investigate task-specific effects in the context of perceptual decision making [361, 86].

Second, owing to the potential variability of online testing environments between participants, we designed our experiment to be completed in one testing session. Perceptual learning studies showed that behavioural performance of coherent motion discrimination improves steadily over multiple testing sessions across several days [470, 246]. It would be of interest to examine if repetitive training modulates the behavioural change between single and multiple information sources.

In conclusion, when sensory information is separated into independent apertures,

perceptual decisions are less accurate. Our cognitive and neural-mass modelling showed two selective neurocognitive mechanisms underlying the behavioural effect, a change in the signal-to-noise ratio of the accumulation process and the speed-accuracy trade-off, depending on the congruency of multiple sensory sources. These findings suggest that both attentional demands and endogenous response strategies influence flexible decision making in humans.

Chapter 5

Spiking neural network model: Inferential decision making

This chapter is based on the work published in the Proceedings of the *42nd Cognitive Science Annual Conference* [103] in collaboration with Peter Duggins (conceptualisation, methodology, software), Prof. Chris Eliasmith (methodology, supervision of modelling), Dr Szymon Wichary (conceptualisation, data collection, interpretation of the results, supervision). The author's contribution to this work include methodology, software implementation, modelling, parameters hyper-optimisation and writing.

Decision-making requires the coordination of anatomically and functionally distinct cortical and subcortical areas. While previous computational models have studied these sub-systems in isolation, few models explore how decision holistically arises from their interaction. The model of inferential decision-making presented in this Chapter exemplifies the application of NEF. NEF offers a systematic method for performing hypothesis testing in the context of neurally realistic models [113]. Here, we propose a spiking neuron model that unifies various components of decision-making. We show that the model performs an inferential decision task in a human-like manner, exhibiting two typical strategies of choice.

5.1 Background

In inferential decision-making (IDM), an agent must gather evidence about the value of possible actions then choose between them. IDM tasks pervade natural and artificial environments, presenting a unique set of cognitive challenges related to the acquisition and processing of information [309]. Notably, these tasks are dynamic and stochastic: observing and processing sensory data takes time, and data are often sampled randomly from underlying probability distributions.

To effectively solve IDM tasks like appraising a potential mate or diagnosing an illness, animals must deploy cognitive processes that deal with time and uncertainty [443]; they must accumulate the probabilistic evidence for each potential choice, continually compare those choices until some decision criteria are reached, and finally implement the preferred option. Although most humans perform these steps intuitively, the underlying cognitive operations are far from trivial: an agent must internally represent the sampled evidence, judge the quality of that information, remember evidence for multiple choices, and choose when to make a decision. All of this must be done flexibly to meet task demands and account for the environmental context [443].

We are interested in the neural and cognitive processes that underlie decision making on IDM tasks, specifically how brains manage trade-offs between speed and accuracy. It makes evolutionary sense that, in environments where an agent must move towards desirable outcomes expediently, IDM should proceed quickly and accurately [125]. Unfortunately, either speed (time taken before making a choice) or accuracy (about the value/validity of the choice) is often sacrificed in dynamic and uncertain environments. We believe that brains use biasing to flexibly adapt IDM processes to favour either speed or accuracy, and that individuals may adopt different strategies, either contextually or habitually, on these tasks [443, 444]. This problem is heightened by the limited capacity of cognitive systems running on physical hardware that limits the precision of numerical representations and the speed of information transfer.

Consider an IDM task in which an interviewer must decide which of four job candidates would be the most productive worker. The interviewer may call a candidate into her office, ask them a question, and numerically score their response. By repeatedly

calling different candidates and asking different questions, the interviewer dynamically collects evidence about individuals' potential productivity. Each of the interviewer's questions probes one of several *attributes*, such as intelligence, initiative, or creativity. Because the interview process is imperfect, the score assigned by the interviewer to, say, a candidate's intelligence may not reflect their true intelligence: the *value* assigned to that attribute is sampled stochastically. Each attribute has an associated *validity*, which indicates the partial information it conveys about productivity and the confidence that the scoring is accurate. To decide which candidate is the most productive, the interviewer repeatedly measures one candidate's value for one attribute, multiplies this score by the attribute's validity, and integrates the result with existing evidence for the candidate. Once the interviewer has asked each candidate each question, she can select the individual with the greatest evidence; assuming that the validities are appropriate and the interviewer can do arithmetic, this process will lead to the correct selection. Although this algorithm is effective, it may be unnecessarily slow; for example, if one candidate receives a perfect score, there is no need to continue interviewing. Various heuristics could similarly improve speed, but many come at the expense of no longer guaranteeing the correct selection [443, 375]. Our goal is to study how the cognitive biases that realise these heuristics affect IDM in humans.

In this study, we present an anatomically mapped, spiking neuron model of the cognitive processes underlying IDM. The model performs a simplified version of the task for which human behavioural data is available [444]. We show that manipulating model parameters governing the dynamic threshold for decision making leads to heterogeneous agent behaviour and reproduces a diversity of strategies used by humans on the task. Specifically, the distribution of cues requested before making a decision is comparable in model agents and human participants. Furthermore, we observe losses in accuracy associated with faster decisions at similar levels in the model and in humans. We conclude by discussing extensions of the model to expand its cognitive realism and to account for emotional bias.

5.1.1 Cognition

Inferential decision making with multiple cues requires many interrelated cognitive operations. Normative theories of choice rooted in economics postulate that before making a decision, humans should process all available information and carefully weigh the cues: that is, multiply validities with the cue values and add the results [52, 125]. Descriptive theories of choice, on the other hand, postulate that people frequently use heuristics to simplify decision problems [141]. Although generally successful, heuristics trade off choice accuracy for speed, particularly in situations that require integrating many pieces of information. Recent studies summarised these approaches in a unified framework called *Bottom-Up Model of Strategy Selection* [443]. The model assumes that the use of the rational *Weighted Additive* strategy and the heuristic *Take The Best* can be explained by one unifying, neurophysiologically plausible mechanism. According to this hypothesis, there are three processes that form the bottom-up mechanism of decision strategy selection and lead to the final choice: cue weight computation, gain modulation, and also weighted additive evaluation of alternatives. Overall, many other heuristics were proposed as plausible models of choice, particularly under time pressure [335] and stress [444]. Current evidence suggests that these factors bias decision making, such that choices are made earlier and with less information, in order to adapt to urgent or demanding situations.

5.1.2 Neuroanatomy

Decision strategies are composed of distinct cognitive operations associated with activity in distinct brain structures. In accordance with growing literature, we assume that the following areas are responsible for the various operations required by our computational model. Ventromedial prefrontal cortex, orbitofrontal cortex (OFC), and dorsolateral prefrontal cortex (dlPFC) perform cue evaluation (weighting) [317]. dlPFC is additionally responsible for temporary retention of cue weights in working memory and performs weighted evidence accumulation [94]. Option selection is executed by the basal ganglia (BG) and its recurrent connections to cortex [48]. Within these loops, presupplementary motor area (pSMA) and the right inferior frontal cortex

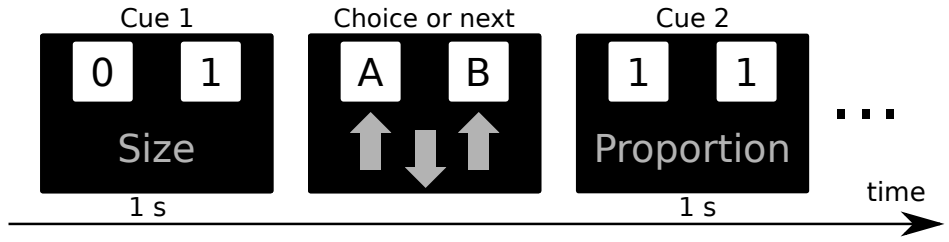

Figure 5.1: *Demonstration of the experimental paradigm. Every cue is presented for 1 s, then the choice is made participants respond by pressing one of three buttons, indicating their selection of an option **A**, **B**, or request for more information.*

(rIFC) act as gates to the striatal input nuclei of the basal ganglia [129], allowing action selection only once an evidence threshold has been passed [12]. Finally, biases of predecisional information processing due to time pressure or stress can arise through the action of the brainstem locus coeruleus norepinephrine system (LC-NE) and its bidirectional connections with the prefrontal cortex, particularly the anterior cingulate cortex (ACC) and OFC [14]. LC neurons exhibit two modes of activity: phasic and tonic. Phasic LC–NE modulation was shown to affect neuronal dynamics and to enhance performance in terms of reward rate in perceptual decision making [106, 107]. LC neurons exhibit also a tonic activity mode associated with disengagement from the current task and exploration [14]. Empirical results generally support the role of all aforementioned these structures in complex decision making with fast choices and little information [291].

5.2 Experimental task

The participants of this task were 17 adult volunteers (9 women, age (years) M = 23, SD = 2.54). Participants provided written informed consent in accordance with the Declaration of Helsinki, under a protocol approved by the Ethics Committee of the Faculty of Psychology, University of Warsaw.

In our IDM task [445], participants begin by memorising validities associated with six cue attributes. Specifically, a computerised probabilistic inference task consisted of making decisions about, which of two diamonds was more expensive. The diamonds

were described with the following cues: size, overall proportions of the diamond, crown proportions, pavilion proportions, size of table and colour. The validities in this task, which range from 0.706 to 0.600, are compensatory: they do not differ significantly from one another, encouraging participants to pay attention to all cue information rather than (heuristically) considering only the top few attributes. During the task, participants are simultaneously shown the values (0 or 1) of objects A and B for one attribute (see Figure 5.1). The participants respond by pressing one of three buttons, indicating their selection of A, selection of B, or a request for more information. If they choose the latter, the current display is replaced by another pair of values for the next attribute. This is repeated until the participant makes a choice or until all attributes have been exhausted, at which point a choice is forced. Attributes are presented in order of highest-to-lowest validity. Participants are not provided feedback on whether their final choice was correct. This task is repeated 48 times; behaviour on each trial is quantified by the number of cues requested before the final decision and whether the choice corresponded to the highest-value option.

5.3 Spiking neural network model

5.3.1 Neural Engineering Framework

The NEF [114] describes how spiking neural activity may represent a time-varying, vector-valued signal $\mathbf{x}(t)$ such as value, validity, or evidence. A neuron spikes most frequently when presented with its particular 'preferred stimulus' and responds less strongly to increasingly dissimilar stimuli (i.e. values of $\mathbf{x}(t)$). In the NEF, each neuron i is accordingly assigned a preferred direction vector, or *encoder*, $\mathbf{e}_i$ [113]. To produce a variety of tuning curves and firing rates that match electrophysiological variance within the brain, each neuron is also assigned a unique gain α_i and bias β_i. These quantities determine how strongly an incident vector $\mathbf{x}(t)$ drives the neuron:

$$I_{in}(t) = \alpha_i * (\mathbf{e_i} \cdot \mathbf{x}(t)) + \beta_i \tag{5.1}$$

where $I_{in}(t)$ is the current flowing into the neuron and $(\cdot)$ is the dot product between the encoder and input vector. So long as there is a well-defined relationship between input current and resulting firing rate, the neuron's activity can be said to encode the vector $\mathbf{x}(t)$. A distributed encoding extends this notion: if $\mathbf{x}(t)$ is fed into multiple neurons, each with a unique tuning curve defined by $\mathbf{e}$, α, and β, then each neuron will respond with a unique spiking pattern $a_i(t)$, and the collection of all neural activities will robustly encode the signal.

Spike generated by a neuron are smoothed using a post-synaptic filter $h(t)$ [113], an operation that approximates the biophysical processes in which incident spikes are translated to post-synaptic current

$$a_i(t) = \sum_s h(t) * \delta_i(t - t_s) \tag{5.2}$$

where the summation is over the neuron's spikes s, each represented as a delta-function $\delta_i(t - t_s)$, and $(*)$ denotes convolution with the low-pass filter.

For neural encoding to be meaningful, there must be methods to recover, or decode, the original vector from the neurons' activities; together, encoding and decoding constitute neural *representation*. The NEF identifies neural *decoders* $\mathbf{d}_i$ that either perform this recovery or compute arbitrary functions, $f(\mathbf{x})$, of the represented vector. A functional decoding with $\mathbf{d}_i^f$ allows networks of neurons to *transform* the signal into a new state, which is essential for performing operations such as value-validity multiplication. To compute these transformations, a linear decoding is applied to the activities of the neural population:

$$\hat{f}(\mathbf{x}(t)) = \sum_{i=0}^{n} a_i(t) * \mathbf{d}_i^f, \tag{5.3}$$

where n is the number of neurons and the hat notation indicates that the computed function is an estimate. To find decoders $\mathbf{d}_i^f$ that compute the desired function, we use least-squares optimisation to minimise the error between the target value $f(\mathbf{x}(t))$ and the decoded estimate $\hat{f}(\mathbf{x}(t))$ [113]. The general-purpose Nengo neural simulator [33] includes methods to optimise these decoders for the specified transformations. Connection *weights* between each presynaptic neuron i and each postsynaptic neuron j

combine encoders and decoders into a single value used during simulation:

$$w_{ij} = \alpha_j \mathbf{e}_j \cdot \mathbf{d}_i^f. \tag{5.4}$$

Finally, the NEF specifies methods to build neural networks that implement any dynamical system, including linear systems of the form $\dot{\mathbf{x}}(t) = A\mathbf{x}(t) + B\mathbf{u}(t)$. To do so, the matrices A and B must be modified to account for the dynamics that naturally occur when using neurons with non-instantaneous synapses. Nengo performs this optimisation for the specified target dynamics; this is essential for constructing networks that include the recurrent connections required for working memory, evidence accumulation, and choice competition [114, 113].

5.3.2 Details of the model

Simulated spiking neuron activity was generated with the Nengo simulator [32]. For our simulations, we used a simplified leaky integrate-and-fire (LIF) model of a neuron to reduce the computation time of the simulation [193]. The subthreshold evolution of LIF neruon voltage V is described by

$$\frac{dV(t)}{dt} = -\frac{1}{\tau_m}(V - I(t)R)$$

where τ_m is the membrane time constant, R is a passive membrane resistance and $I(t)$ is the input current. When the membrane crosses a threshold V_{thr} a spike is emitted and cell is reset to its resting state V_{res}. Additionally refractory time constant τ_{ref} (where a neuron is locked in resting state after emitting a spike) is set to 4 ms. All the spikes emitted by a neuronal ensemble to a different region are filtered with a 100 ms post-synaptic filter (apart from **rIFC** to **pSMA** inhibitory connections with faster 10 ms post-synaptic time constant). Tuning curves controlled by V_{thr} and τ_m were fitted using NEF (see Section 5.3.1).

The anatomy mapping of the model architecture is summarised in Figure 5.2 and outlined here, with detailed descriptions of key components in the following paragraphs.

A two-dimensional vector representing the values of A and B for the currently

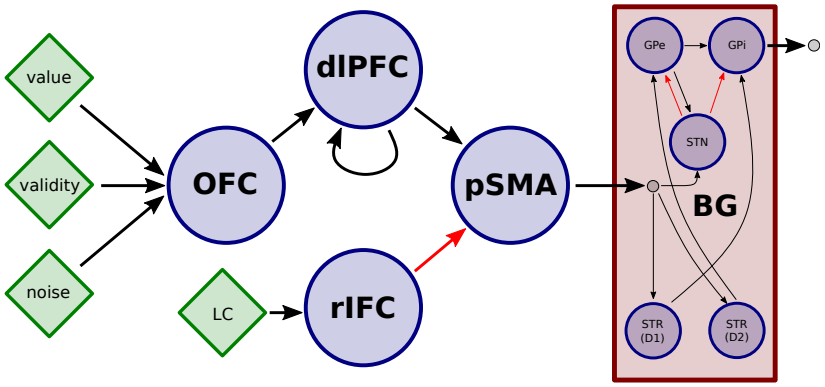

Figure 5.2: *Model schematic. Green diamonds are external inputs, blue circles are spiking neuron populations, and the red box is the **BG** network. Red connections are inhibitory. See text for details on represented quantities and cognitive operations.*

displayed attribute is provided as external input to a population labelled **OFC**. This population also receives the 'remembered' cue validities; we choose to model this recall process as a noisy perturbation of externally-supplied validities*. Neurons in **OFC** thus represent both the perceived cue value and the remembered validity for the current attribute. Connection weights between **OFC** and **dlPFC** multiply values by validities and send the result to the two-dimensional **dlPFC** population. Recurrent connections within **dlPFC** implement integration, leading to the accumulation of evidence from **OFC** as additional cues are presented. When the difference between accumulated evidence for A and B exceeds a dynamic threshold, neurons in **rIFC** disinhibit the **pSMA** population. This allows information to flow from **dlPFC** to the **BG** network, where mutual inhibitory competition selects the option with the greatest evidence as a final output. If **BG** does not output a selection after one second of value/validity input, then the next cue is presented for one second, and so on.

The **dlPFC** population is a neural *integrator*, a system which maintains its currently represented value while additively incorporating any inputs. This network has previously been used in neural models of working memory, where it has reproduced activity in PFC and behaviour on several working memory tasks [113]. The system is described by the target dynamics $\dot{\mathbf{x}} = B\mathbf{u}$; notice that changes in the represented

*Imperfect recall was realised by adding a smoothed 10 Hz bandpass-limited white noise signal with RMS = 0.05

value $\dot{\mathbf{x}}$ do not depend on the represented value $\mathbf{x}$ itself, but only on the input $\mathbf{u}$. This implies that the integrator will remember the current evidence perfectly, and steadily add any input evidence to arrive at a new value. However, because the integrator is implemented in noisy spiking neurons, the feedforward and recurrent connections do not perfectly implement these dynamics: added evidence is slightly distorted, and memory will decay given enough time. The `evidence` population includes 2000 LIF neurons with firing rates up to 400 Hz; see Figure 5.3 in Section 5.4 for an accuracy comparison to optimal integration [71, 368].

The **rIFC** population determines when the accumulated evidence (E) is sufficient to make a decision. Functional connections between **dlPFC** and **rIFC** deliver the input

$$\mathbf{u}_{\text{evidence}}(t) = |E_A - E_B|, \tag{5.5}$$

the absolute difference in accumulated evidence for A vs B, into the latter population. Neurons in **rIFC** are initialised such that all tuning curves have (a) negative slope (larger inputs $\mathbf{u}$ produces less firing) and (b) intercepts fixed at T_{int} (inputs larger than $\mathbf{u} = T_{\text{int}}$ are rectified, producing zero activity). Together, these constraints ensure that neural activities remain positive for $\mathbf{u} < T_{\text{int}}$ and go silent for $\mathbf{u} > T_{\text{int}}$. **rIFC** connects to **pSMA** with strong inhibitory connections[†], such that any activity in **rIFC** dampens all activity in **pSMA**, restricting the flow of information to **BG**. When $\mathbf{u}$ surpasses the threshold, disinhibition lifts this gating, and a decision follows shortly thereafter. Large T_{int} thus implies that the difference in accumulated evidence for A and B must be large before a selection is made.

To simulate time pressures and heuristics that favour speed over accuracy, we model an additional input $\mathbf{u}_{\text{decay}}$ to **rIFC**. This input is additive with $\mathbf{u}_{\text{evidence}}$ and grows linearly during the course of the trial

$$\mathbf{u}_{\text{decay}}(t) = T_{\text{decay}} * t, \tag{5.6}$$

effectively decreasing T_{int} as more cues are presented. Large T_{decay} thus implies that

[†]The **pSMA** population is a gated communication channel that performs no transformations on its own.

the decision threshold shrinks more quickly as time progresses. $\mathbf{u}_{\text{decay}}$ is also set to a large positive value when $t > 5.5$ s, which removes gating at the end of the trial and forces a selection.

The **BG** population is based off an anatomical reconstruction the basal ganglia and implements winner-take-all competition between action alternatives [380]. As with the neural integrator, this network has been used in numerous functional brain models as part of the action selection system [113]. Here, it is simply used to select the action with the greater evidence value, effectively amplifying differences between E_A and E_B to produce a clear distinction for downstream systems. We sample the outputs of **BG** and mathematically determine which of the three action options (choose A, choose B, request more) is greater, assuming their difference surpasses a noise threshold of $T_{\text{noise}} = 0.1$ [379].

Each model agent is initialised with unique T_{int}, T_{decay}, and random seed for generating neurons parameters $\mathbf{e}$, α, and β. As with human participants, model agents repeat the task 48 times, and the number of cues and correctness on each trial are recorded.

5.3.3 Open data and scripts

We have made all analysis scripts and behavioural data publicly available: https://github.com/psipeter/decision_strategies.

5.4 Simulation results

To clarify the dynamics of the model and investigate the causes of correct and incorrect choices, we begin by looking at time series for the state variables represented in neural populations. Figure 5.3 shows three trials from different agents, including one successful choice and two failures. In the first trial (top), choice B has positive values for the first three cues, while choice A has zero value. The difference in accumulated evidence steadily grows (blue vs. red line) until the dynamic decision threshold is reached around $t = 2$ s. According to [14], the relationship between LC-NE tonic activity

and phasic gain takes the form of an inverted 'U'. When the LC-NE tonic activity is low, LC-NE phasic response to cues is weak (gain increase is small). When the tonic LC activity is moderate, its phasic response is strong (gain increase is high), and again when the tonic activity is high, the phasic response is low (gain increase is small). The current models postulate the following: when the initial tonic LC activity is low, then ACC drive will increase it to the moderate level, which will then lead to increased phasic responses to incoming decision cues. This will also lead to better gating of these cues into the rIFC and, on the psychological level, their higher impact on the option evaluation [443]. In our simplified model, from $t = 2$ s an inhibition from **LC** begins quieting the tonically active threshold population **rIFC** (black line). Once this activity decays to zero ($t = 2.5$ s), the **pSMA** gate (dot-dash line) becomes disinhibited, and information flows from **dlPFC** to **BG**, which quickly selects the dominant choice. At the end of the third cue ($t = 3$ s), an external check registers that **BG** has made a selection, and the agent is said to have chosen after three cues. A demonstration of this choice strategy is presented on the simulation video clip: https://www.youtube.com/watch?v=Ksm1hZbzBbs.

In the second trial (middle), both A and B have positive values, making the two choices barely distinguishable. The difference in accumulated evidence never exceeds the agent's large threshold. At $t = 5.5$ s, **rIFC** is externally inhibited, opening the gate and effectively forcing a decision. By this point, noise-induced errors have accumulated in the **dlPFC** representation; when this evidence is fed to the **BG**, the agent incorrectly estimates that A's evidence exceeds B's, and it selects the wrong choice. In the third trial, the evidence initially favours B; because this agent has a small threshold, its decision criteria is met by $t = 3$ s, leading to selection of B. However, the remaining evidence from cues four through six tip the balance in favour of A, making the agent's early choice ultimately incorrect. A demonstration of this 'impatient' choice strategy is presented on the following video clip: https://www.youtube.com/watch?v=Rckz5sEIoPM.

On the third trial, the agent's small threshold is exceeded during the third cue, leading to selection of B; however, the remaining evidence from cues four through six tip the balance in favor of A, making the agent's choice ultimately incorrect. By this point, small noise-induced discrepancies between representations in evidence (solid line)

and the mathematically computed sum (dashed line) have accumulated, so the agent incorrectly estimates that A's evidence exceeds B's, and chooses the wrong option.

What is an appropriate decision threshold to ensure both accuracy and speed? To answer this question, we examine individuals (both humans and agents) who strongly favour either speed or accuracy. We initialised a population of agents with random default thresholds, threshold decays, and neuron parameters, then had them perform 48 trials of our IDM task. Figure 5.4 compares the behaviour of fast, inaccurate decision makers and slow, accurate decision makers; behaviour is quantified by plotting the number of cues inspected before a decision across all trials. As expected, an agent with small default threshold ($T_{\text{int}} = 1.41$) and large threshold decay ($T_{\text{decay}} = 0.41$) typically makes selections after two or three cues, but has lower accuracy (77%) than a large default, small decay agent ($T_{\text{int}} = 3.00$, $T_{\text{decay}} = 0.33$), who views four or more cues before making a choice (accuracy 94%). The behavioural distributions and accuracies of both agents are closely aligned with relevant human participants.

The relationship between speed and accuracy is readily apparent in Figure 5.5, which plots individuals' mean accuracies as a function of their mean cue requests. The trend is clear and intuitive: spending more time gathering information and performing active inference leads to more correct choices. The highest accuracies are obtained by individuals who wait for five or six cues before making a choice, although performance rarely surpasses 90% due to the difficulty of remembering and summing non-compensatory cues to high accuracy. Still, even the hastiest decision makers, who choose before the fourth cue, have accuracies above 70%, making these strategies a viable alternative when faced with strong time pressures. As before, trends in the simulated data agree with the empirical data.

5.5 Discussion

In this research we use a spiking neuron model for simulating the two most common decision strategies in the IDM task [443]. We take advantage of the NEF - a mathematical theory of how biological neural systems can implement a wide range of dynamic functions. Previously, NEF has been applied to both low-level systems (e.g. involved

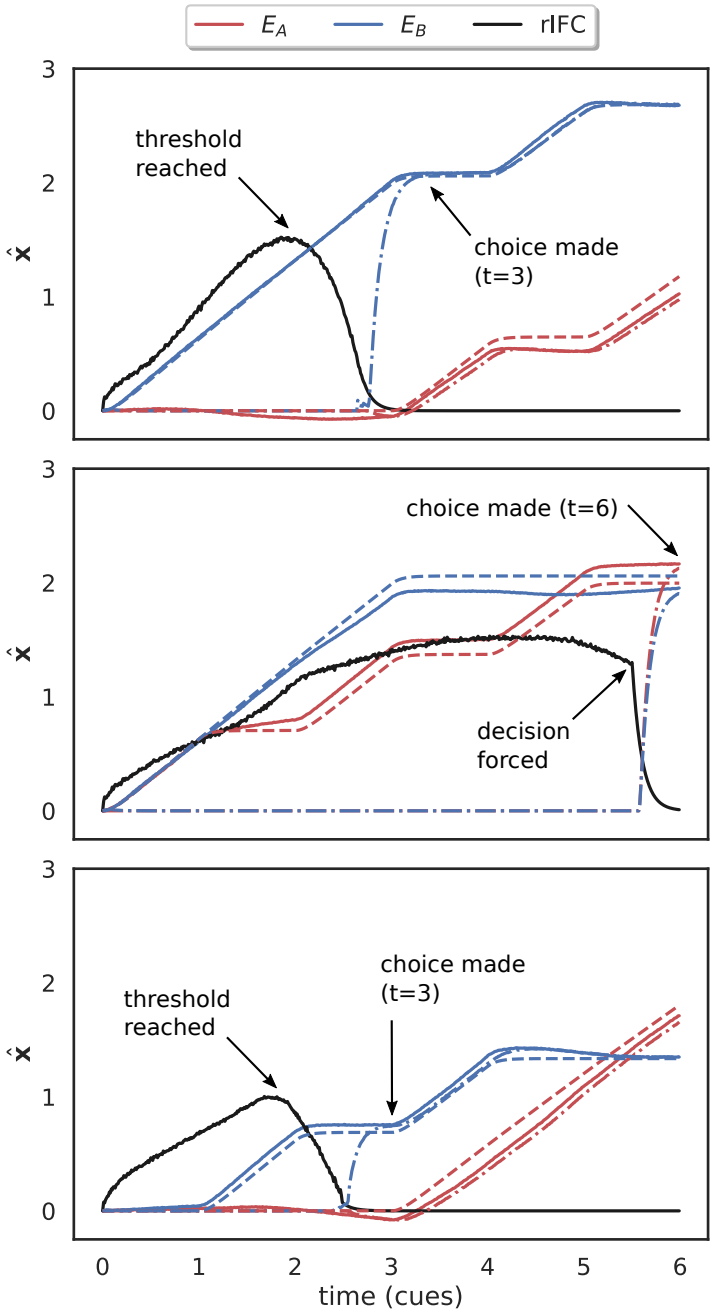

Figure 5.3: *Time series of neural representations* $\hat{\mathbf{x}}(t)$ *in three model agents. Solid lines represent accumulated evidence computed by* ***dlPFC*** *neurons in the model, while dashed lines represent the ideal accumulation (which means a perfect accumulation of the evidence from weighted cues). Dot-dash lines represent information flowing through the* ***pSMA*** *gate, which opens when the thresholding population* ***rIFC*** *(black line) has been inhibited.*

in stabilising eye position), but also higher-level cognitive phenomena (e.g. working memory) [113, 44, 33]. For the purpose of this task, several different modelling frameworks could have been used. For example, the `Brian2` Python package offers an easy programming interface for creating models of spiking neurons [382]. Due to the closer level of biological realism, the use of `Brian2` for large networks (in our case: 7200 neurons) leads to longer simulation times. This could be reduced, however, by using mean-field approach. Another modelling option was to use the DCM to infer the causal architecture of distributed dynamical systems. DCM was developed for estimating coupling among macroscopic brain regions based on the comparison with the functional data (i.e. fMRI, or M/EEG) [132]. In our study, however, we compare the model performance only with behavioural data. Furthermore, NEF allows for better between levels analysis (from spiking activity of cells within specific subcortical regions to an inferential decision) giving a room for an additional experimental verification with extracellular recordings of neuronal activity.

Our IDM model was designed to recreate the anatomy and cognitive function of the human decision making system using populations of spiking neurons. Although the hu-

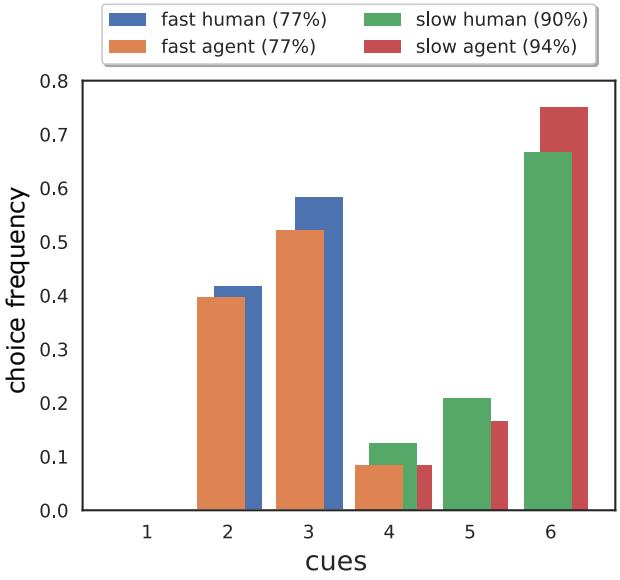

Figure 5.4: *Behavioural distributions of steps needed to make a decision for humans and agents.*

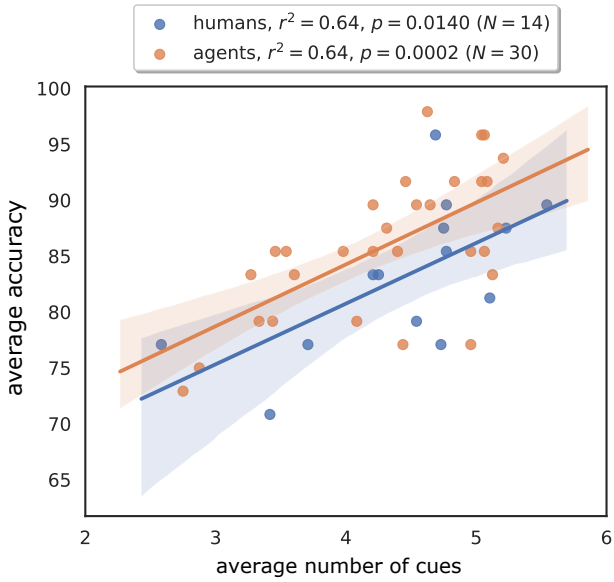

Figure 5.5: *Average decision accuracy as a function of average cues requested.*

man decision making system is too complex to capture in a simple network, we believe that our computational model distils many of its core features into a functionally-clear circuit. **OFC** receives inputs from sensory and memory systems and passes it to **dlPFC**, which performs multiplication. The result is incorporated into actively maintained representations of the evidence for various choice alternatives; such a working memory is also realised in **dlPFC**. Connections between **dlPFC** and **BG** are responsible for forwarding the utility of action alternatives to an action selection system, but these connections are gated by intermediate cortical-subcortical structures like the **pSMA**, ensuring that the individual waits until an appropriate moment to enact a behaviour. Only once a controllable threshold, represented in cortical areas like the **rIFC**, has been exceeded will **BG** be freed to make a final decision and implement a motor response. This thresholding may be affected by subcortical modulation, including noradrenergic projections from **LC**, which assess environmental context, time pressures, and the individual's emotional state before deciding whether to bias decision making towards speed or accuracy. Although the particular functions computed on neural connections in our model are specific to this IDM task, we believe that this architecture

may be utilised for a variety of decision making tasks involving dynamic inference on stochastic evidence.

Even though anatomical realism was key in this research, for simplicity, our model omits several important anatomical structures. This is most notably the dopaminergic system, which is implicated in coding of positive utility, reward, and arousal, as well as in working memory processes [359]. Another important region that has been ignored in this modelling is the nucleus accumbens, which is associated with reward mechanism and stimuli reinforcement circuits [190].

We showed that our model agents reproduced two measures of human behaviour: the number of cues examined before making a decision, and the relationship between accuracy and number of examined cues. We found that agents with smaller dynamic thresholds chose quickly but inaccurately, while agents with larger thresholds chose slowly but with higher accuracy. The distributions of simulated choices among model agents matched both extremes of human behaviour, and the trendline of speed versus accuracy across a variety of agents matched behaviours from a small human dataset. These successes suggest that the model may capture key functional aspects of the decision making apparatus in humans, especially how biases towards speed or accuracy, either situational or ingrained, may change behaviour.

Future work can profitably proceed in several directions. To increase the model's cognitive realism, an associative memory system could be trained by pairing 'keys' for the input attributes (unique high-dimensional vectors representing each verbally-defined attribute) with sensory inputs describing the cues' validities (one-dimensional values). This pre-training would utilise an encoder-based learning rule that has previously been used in `Nengo` models of associative memory [422]. After training, presentation of the cue would elicit noisy recall of the associated validities, which would then be routed to **OFC** as in the current model. We are also interested in extending the model to more complex IDM tasks, including the 'job interview' task described in the introduction. This task involves more choice alternatives and a wider range of cue validities; it also requires participants to view values/validities for one option at a time, and allows them to chose which option/attribute to query. This freedom introduces an extra dimension of exploration, in which the decision about which 'questions to ask'

may interact in interesting ways with the current evidence for various options. Finally, experimental data is available for a variant of this task in which participants were shown neutral or aversive images before performing the task, priming them with an emotional state that affected their decision making [444]. We would like to investigate whether shifts in the dynamic threshold can capture this emotional biasing, specifically whether valenced arousal states bias individuals towards faster, less accurate decision processes.

Chapter 6

Discussion

All models are wrong, but some are useful.

George Box

6.1 Summary of the results

In this thesis, I presented four examples of how models and computer simulations contribute to generating and testing hypotheses at various levels of the brain hierarchy: from macroscopic resting-state brain networks to the spiking activity of neural masses. The key findings are summarised below.

In the experimentation described in Chapter 2, the behavioural data (accuracy and response times) and EEG signals were recorded while participants performed a voluntary decision-making task. Classical decision theories predict longer response times for settings where the outcome of a binary decision is the same [325, 391]. This contrasts, however, with the notion of an optimal decision maker who should make fast and random choices. In our experiment, we observed that a higher reward probability accelerated the responses. We also found an effect of a preference towards one option, which helps solve the ambiguous choice problem under time pressure [465]. By using hierarchical Bayesian parameter estimation for an accumulator model, we showed that the probability of reward and preference were independently associated with changes in

the speed of accumulation of evidence, but not with visual encoding or motor execution latencies [198]. EEG-informed modelling revealed that the rates of changes between N100 and P300 ERPs are modulated by the accumulation rates on a trial-by-trial basis. Overall, our findings suggest that reward probability and spontaneous preference collectively shape voluntary decisions between equal options and provide a mechanism to prevent indecision or random behaviour [465].

In Chapter 3, a different kind of model was used. We characterised the statistical regularities among selected resting-state networks and compared them between patients suffering from JME and a healthy control group. We used the energy landscape method for macroscopic brain signals [40, 437, 389]. The pMEM was fitted to the instantaneous MEG oscillatory power. It provided a good estimation of the occurrence probability of network states from only 6 minutes of recordings, due to the high sampling rate of MEG [233]. Then, we used the energy values derived from the pMEM to construct the energy landscape, describing the probability of occurrence of transitions between the states. The landscapes of JME patients displayed fewer local energy minima than the landscapes of the controls, and had elevated energy values for the FPN within the theta, beta, and gamma frequency bands. Furthermore, simulations of state transitions showed that the proportion of time the FPN was occupied within the basins of energy minima was shortened for JME patients. These network alterations were underscored by significant classification of individual participants employing energy values as multivariate features. All in all, our analyses suggested that JME patients had altered multistability in selective functional networks and frequency bands in the fronto-parietal cortices (i.e. theta FPN) [233]. Furthermore, they confirmed that the pMEM is a descriptive, generative, and predictive model for characterising atypical network properties in the brain [437, 389].

In Chapter 4, we combined the experimental work with cognitive modelling, which led to the generation of a hypothesis based on an NMM of decision making. A typical perceptual decision-making experiment involves choosing one of two options from only single source of evidence (e.g. a single stream of dots) [362]. In reality, decisions involve combining the evidence from multiple information sources. Thus, the effect of cognitively linking that information has not been yet thoroughly investigated. In this

project, we designed a new experiment based on the classic random-dot motion discrimination task [54] and recorded the behavioural responses (accuracy of choice and reaction times) using an on-line platform for behavioural experiments. We observed imperfect integration of information from two information sources in both the cases of congruent and incongruent evidence displayed on the screen. We propose that this can be explained by means of changes in the drift rate and threshold of hierarchical DDM. Based on the cognitive model, we hypothesised about the neuronal implementation, suggesting a modification of a previously reported biologically-plausible NMM of decision making under time constraints [454]. Our new NMM reproduced all four conditions of the task.

Finally, in Chapter 5, we proposed the spiking neural network model, informed by brain anatomy, for inferential decision making. Previous computational models have studied the coordination of functionally distinct cortical and subcortical areas in isolation [335]. Only a few models have explored how decisions arise holistically from the interaction of the areas of the brain [380]. We proposed a model that unifies various components of the decision-making system and showed that it performs an inferential decision task in a human-like manner [444, 103]. Our model includes neuronal populations corresponding to the dorsolateral prefrontal cortex, orbitofrontal cortex, right inferiorfrontal cortex, pre-supplementary motor area, and basal ganglia. It consists of 8000 leaky-integrate-and-fire neurons making in all around 7 million connections. The model realises dedicated cognitive operations such as weighted valuation of inputs, accumulation of evidence for multiple choice alternatives, competition between potential actions, dynamic thresholding of behaviour, and urgency-mediated modulation. We proved that the model was able to reproduce reaction time distributions and speed–accuracy trade-offs similarly to the human subjects performing the task [103].

6.2 Limitations and future directions

Here, I briefly summarise general limitations and future directions of the work discussed in this thesis. The specific comments regarding particular projects are included in the respective Chapters (2, 3, 4, and 5).

Neuronal and cognitive models have a rich variety of hierarchies and parameters. In this thesis, I presented only a small fraction of all the models that are used in neuroscience. For example, Bayesian modelling has gained a lot of attention in cognitive science recently [212, 415]. Bayesian models rely on the assumption that the brain approximates statistically optimal solutions of the performed tasks. More formally, they can be viewed as generative models of the joint distribution $p(D, I)$ (where D denotes the data represented in the brain, and I the sensory input). The joint distribution is equal to the product of the *prior* $p(I)$ over all possible configurations of the input and the *likelihood* $p(D|I)$, which is the probability that the data occurs given the sensory input [415]. Another important line of work is a framework for specifying, fitting and comparing models: *Dynamic Causal Modelling* (DCM) [132]. DCM is designed to investigate the influence between different areas of the brain using multivariate time series (like M/EEG or fMRI) [272]. The approach consists of fitting various models to the time-varying data, and selecting one of the models with a Bayesian model comparison. This thesis, however, does not include any work related to Bayesian brain modelling or DCM, although together they make up an important part of contemporary modelling in neuroscience.

In this thesis, most of the models have focused on isolated components of neural activity. In Chapter 2, we showed the EEG-informed DDM improves the predictive power in voluntary decision task during a deadlock. We used the slope between N100 and P300 ERP components from the central part of the scalp [465]. However, the averaged EEG signal is much richer and the impact of other ERPs on evidence accumulation is not yet known. In Chapter 3, we used a macroscopic model which considered the brain as a graph with 90 nodes (after source reconstruction and parcellation). Yet, due to the numerical limitations of the energy landscape method, we investigated only differences between selected resting state networks: DMN, FPN and SMN [233]. In Chapter 4, we focused on the LIP area of the brain, which has been suggested as being engaged in perceptual decision-making [364, 362]. Nonetheless, in Chapter 5, we provided a more inclusive view of the decision system in the brain. Still, the model is far from complete. Some brain structures were omitted (e.g. nucleus accumbens), or simplified to improve the interpretation of the model with respect to the task [103]. The possible

extension of the study from Chapter 2 could validate the formulated hypotheses by controlling for memory effects originating from other brain areas, not considered in this study [253]. The effects of brain areas ignored in the model from Chapter 5 should be further confirmed by experimental work, e.g. fMRI experiments.

Furthermore, all the projects in this thesis focused on selected aspects of behavioural or electrophysiological recordings. Multi-modal research can give greater mechanistic insight than uni-modal approaches [3, 150]. For example, in perceptual decision making, fMRI can help with localising the brain regions engaged in the decision process [203], whereas M/EEG recordings can help to understand the role of timing in fast-paced evidence accumulation [288, 289]. Currently, even more studies employ multi-modal research for broadening the perspective of the modelling work [3, 150]. An extension of the energy landscape project from Chapter 3 could investigate how the fast-sampled state dynamics of MEG data is related to the low-sampled but more spatially focused landscapes of fMRI recordings [233, 437].

Chapter 4 and Chapter 5 presented computational models of low-level neuronal dynamics. The modelling results can be used to inform the design of future experiments. For example, we could compare the model outcomes with the local-field potential recordings from an experiment using a similar perceptual decision task. The resulting neuronal recordings could test our model-driven hypothesis of increased spiking activity of spatially non-selective neurons for a condition with two visual streams and incongruent information available (Chapter 4, Figure 4.1). Additionally, we plan to extend the study from Chapter 5 to compare the neuronal activity predicted by the NEF with fMRI recordings during inferential decision making in a similar setup [445, 103]. fMRI recordings offer a window into the activity of deep-brain structures that were hypothesised to be involved in this task, e.g. the striatal input nuclei of the basal ganglia.

6.3 Discussion

The human brain is an immensely complicated system. Neurons produce complex activity within anatomical regions, but at the same time no single region works in

isolation [136]. In Chapter 5, we laid out an attempt to build a comprehensive model of inferential decision making, but even this model omitted some critical brain structures (like the nucleus accumbens, known to be engaged in reward and reinforcement circuits [190]).

The complexity of the human brain calls for a reductionist approach, which studies simplified and isolated models of brain circuits, or selected higher-level cognitive processes [239]. Such models are, however, helpful at directing future experimental work towards specific problems or hypotheses [250]. Although models are a simplification of reality, they aim to capture the primary factors of an underlying system. Finding the right model is a complicated task. It requires striking a balance between the plausibility and the interpretability of the model [239]. For example, a model of an LIP area consisting of 10,000 neurons might be closer to the actual implementation than an abstract DSM with two state variables. Nonetheless, it might be more difficult to understand [454]. The ideal model's task is not merely to reproduce all components of the system. A simplified version (i.e. omitting parts of the system not involved in a task or considered activity) proves useful by capturing the true properties of the system or by making reliable predictions [239].

The initial promises of the Human Brain Project (HBP) led to a lot of controversies in the field of neuroscience. On July 7, 2014, an open letter, signed by around 750 researchers involved in the project, was sent to the European Commission [131]. The neuroscientists pointed out that large-scale simulations make little sense unless constrained by data, or specific research questions. At the initial stage of the project, a detailed 'connectome' of a human brain was not yet known [260, 23]. More importantly, there were no precise biological or cognitive hypotheses formulated for these simulations to test [131]. A recently published overview of the current status of the HBP summarised the reviewed objectives of the project [109]. The researchers highlight the latest developments in large-scale brain simulations, the integration of various data modalities, and define specific cognitive applications for the simulation [109].

Another debatable topic is whether one should use more data-driven or more theory-driven approaches [428]. In the traditional theory-driven scenario, mathematically formulated hypotheses are proposed and falsifiable given observed experimental data

[250, 239]. On the other hand, data-based models, such as machine learning, may help to find unknown relationships in the collected data [428]. Although these two approaches are being pitted against each other, they can be complementary [165]. For instance, by revealing hidden structures in the data, a machine learning model could yield a better understanding of a signal generation process and new NMM models with testable hypotheses could then be formulated [142]. One of the aims of this thesis was to stress the benefits of both data-driven approaches and theory-driven approaches, as presented, for example, in Chapter 3 and Chapter 4.

Due to the recent advances of brain imaging techniques, scientists have begun to create ever-increasing amounts of data. This gives rise to another layer of difficulty, as the new datasets are often hard to manually analyse. Indeed, recent years have seen an exponential growth in the applications of machine-learning to neuroscience [142]. Dimensionality reduction techniques, such as ICA or PCA, help to reduce the complexity of the data [428]. However, the role of machine learning is not limited to multivariate data simplification. Machine learning helps to identify predictive variables, or sets benchmarks for simple models of the brain [142]. Specific models can be used as abstract models of a brain function [142].

Given the variety of existing models, finding a robust and general criterion for the utility of a model is difficult. Several components of a good model can be, nonetheless, formulated. A model should be

- **realistic**: this means that it is based on true observations, or a verified theory;
- **precise**: it should be specific in its predictions and at the right level of the analysis;
- **simple**: it should improve our understanding of a process, so if there is a simpler solution to the problem that gives the same results, one should opt for that one (unless doing this would contradict the two previous points);
- **testable**: it should be possible to verify its predictions via further experiments (see also Figure 1.1).

The last point is probably the most important. This is emphasised by the fact that psychology was reported to have undergone a 'reproducibility crisis', where it was impossible to replicate 50% of the experimental paradigms [122, 450]. The remedy for this problem can be found in different fields. In machine learning, the accuracy of a model is tested on a *validation* set, which is a subset of the data that has not been used for training the model; one must also avoid any hyper-optimisation of the parameters [428, 232]. Thus, experimental work could record not only the data needed to achieve the desired effect size, but also a smaller sample for self-replication [212]. The verification of some models might be in practice more challenging. For example, a model can produce verifiable predictions, but they might be out of reach for existing technology.

Another challenge of modelling in neuroscience is more practical. The unification of the computational modelling tools is a pressing need [144]. Currently, research groups use custom numerical tools, which makes the work difficult to reproduce, compare and verify. Recently, open-source libraries have started to emerge, e.g. Neuron [65], Nengo [113], Brian [382], and The Virtual Brain [347]. The development of crowd-sourced programming projects can benefit the field, as it improves the reliability of the code, helps avoid numerical mistakes, and diversifies the access to the knowledge [110]. For all the projects presented in this thesis, I have accompanied the numerical experiments and analyses with code repositories that contain the computational pipeline and make it easy to re-evaluate the results.

Overall, understating the functioning of the brain begs for developing complementary theories and experimental work. Hence, this requires a closer integration of data-driven bottom–up approaches with more theory-driven top–down studies that start with a behaviour or a cognitive phenomenon to be explained [228]. Only then can the four key components of the ideal model identified in this section be fulfilled. As an example, the emergence of large-scale brain modelling tools [113, 347] could help to integrate NMMs with network approaches to study macroscopic brain behaviour [248]. Additionally, these days, more emphasis is put on improving the reproducibility of computational and experimental work by using pre-registrations and registered reports, publishing the code and data associated with a study, public code reviews, or

replication attempts [255].

6.4 Conclusions

In this thesis, I presented four projects that exemplify the applications of different modelling types at various levels of brain organisation. This is coherent with the contemporary state of the research, where modelling work looks at the brain through the lenses of a particular scale or a cognitive phenomenon. However, brain research is entering a new phase. With the emergence of novel computational tools, as well as new data and theoretical work, a trend for combining different levels of analysis can be noticed. This may allow building comprehensive models that are able to explain cognition with neurobiologically plausible components.

Appendices

Appendix A

Dissemination of Results

A.1 Peer-Reviewed Publications

- Zajkowski W*, **Krzemiński D***, Barone J, Evans LH, Zhang J. (2020) Breaking Deadlocks: Reward Probability and Spontaneous Preference Shape Voluntary Decisions and Electrophysiological Signals in Humans. *Computational Brain & Behavior*. doi:10.1007/s42113-020-00096-6

- Duggins P, **Krzemiński D**, Eliasmith C, Wichary S (2020) A spiking neuron model of inferential decision making: Urgency, uncertainty, and the speed–accuracy tradeoff. In S. Denison., M. Mack, Y. Xu, & B.C. Armstrong (Eds.), *Proceedings of the 42nd Annual Conference of the Cognitive Science Society* (pp. 1891-1897). Toronto, ON, 2020. Cognitive Science Society.

- **Krzemiński D**, Masuda N, Hamandi K, Singh KD, Routley B, Zhang J (2020) Energy landscape of resting magnetoencephalography reveals fronto-parietal network impairments in epilepsy. *Network Neuroscience*. (Cambridge, Mass.), 4(2), 374–396.

- Lopes MA, **Krzemiński D**, Hamandi K, Singh KD, Masuda N, Terry JR, Zhang J (2021) A computational biomarker of juvenile myoclonic epilepsy from resting-state MEG. *Clinical Neurophysiology*. doi:10.1016/j.clinph.2020.12.021

*Authors contributed equally

- Lopes MA, Zhang J, **Krzemiński D**, Hamandi K, Chen Q, Livi L, Masuda N (2020). Recurrence Quantification Analysis of Dynamic Brain Networks. *The European Journal of Neuroscience*. 53(4), 1040–1059.

- **Krzemiński D**, Michelmann S, Treder M, Santamaria L (2019) Classification of P300 Component Using a Riemannian Ensemble Approach. *Proceedings of MEDICON 2019 conference*. Cham: Springer International Publishing. (pp. 1885-1889) Coimbra, Portugal.

A.2 Pre-prints and pre-registered reports

- **Krzemiński D**, Zhang J (2021) Imperfect integration: sensory congruency between multiple sources modulates selective decision making processes. *bioArXiv*; Experiment pre-registered: https://osf.io/4dn65.

- Thomas KA, **Krzemiński D**, Kidziński Ł, Paul R, Rubin EB, Halilaj E, Black MS, Chaudhari A, Gold GE, Delp SL (2021) Open source software for automatic subregional assessment of knee cartilage degradation using quantitative T2 relaxometry and deep learning. *ArXiv*

A.3 Talks and Presentations

- Poster presentation: 'Sensory Congruency Between Multiple Sources Modulates Perceptual Decision-Making' at *10th Symposium on the Biology of Decision-Making* on-line conference, 2021

- Poster presentation: 'U-Net for Automated Segmentation of Knee Cartilage Imaging' at *MLinPL* on-line conference, 2020

- Oral presentation: 'Integration of Information in the Perceptual Decision Making Task' at *Neuromatch 3* on-line conference, 2020

- Poster presentation: 'Selective network dynamic abnormalities in juvenile myoclinic epilepsy revealed by MEG energy landscape' at *Neuroscience* conference in Chicago, 2019

- Poster presentation: 'Supervised text classification for cohort selection in clinical trials' at *Healtac* conference in Cardiff, 2019

- Oral presentation: 'Energy landscapes of spontaneous brain activity of juvenile myoclonic epilepsy patients' at *MEG UK* conference in Cardiff, 2019

- Poster presentation: 'Variability of MEG signals in epilepsy' at *Science Polish Perspectives* conference in Oxford, 2018

- Poster presentation: 'The effects of reward certainty on voluntary choices: an EEG study' at *Organisation for Human Brain Mapping* conference in Singapore, 2018

- Poster presentation: 'Energy landscape analysis of MEG resting state data' at *MEG UK* conference in Londonderry/Derry, 2018

Appendix B

Supplementary Material to Chapter 4

Supplementary methods

The neural-mass model of perceptual decision

We used the two-state neural-mass model [454] in the following form:

$$\frac{dS_i}{dt} = -\frac{S_i}{\tau_S} + (1 - S_i)\gamma r(I_{\text{syn},i}) \ , \tag{B.1}$$

$$r(I_{\text{syn},i}) = \frac{aI_{\text{syn},i} - b}{1 - \exp(-d(aI_{\text{syn},i} - b))} \ , \tag{B.2}$$

$$\begin{cases} I_{\text{syn,L}} = & J_{\text{L,L}}S_{\text{L}} + J_{\text{L,R}}S_{\text{R}} + I_{in,L} + I_{\eta,\text{L}} \\ I_{\text{syn,R}} = & J_{\text{R,R}}S_{\text{R}} + J_{\text{R,L}}S_{\text{L}} + I_{in,R} + I_{\eta,\text{R}} \ , \end{cases} \tag{B.3}$$

$$\tau_\eta \frac{dI_{\eta,i}}{dt} = -I_{\eta,i} + \eta_t\sqrt{\tau_\eta}\sigma_\eta \ , \tag{B.4}$$

where the index $i = \text{L}, \text{R}$ refers to two neural populations selective for leftwards and rightwards choices. The state variable S describes the NMDA gating variable (fraction of open gates). It can be shown that S has a bijection mapping on pre-synaptic firing rates [454]. r describes the population firing rate function that depends on the synaptic input current $I_{\text{syn},i}$. The following values of the parameters were used: $a = 270(\text{VnC})^{-1}$,

$b = 108$ Hz, $d = 0.154$ s, $\gamma = 0.641 \cdot 10^{-3}$, $\tau_S = 100$ ms.

The synaptic input current $I_{\mathrm{syn},i}$ combines recurrent inputs, mutual inputs, external inputs that relate to sensory information ($I_{\mathrm{in},i}$, Equation 1 in the main text) and the noise current $I_{\eta,i}$. The following symmetric synaptic coupling parameters were used: $J_{\mathrm{A,A}} = J_{\mathrm{B,B}} = 0.2601$ nA and $J_{\mathrm{A,B}} = J_{\mathrm{B,A}} = 0.0497$ nA, $J_{\mathrm{ext}} = 5.2 \cdot 10^{-4}$ nA·Hz^{-1}.

The noise current $I_{\eta,i}$ is integrated with $\tau_\eta = 2$ ms (time decay of AMPA receptor activation) and random variable sample from Normal distribution η. The variance of the noise factor is kept constant $\sigma_\eta = 0.01972$.

For simulations, we use standard Euler's integration method with a time step of 1 ms.

Supplementary figures

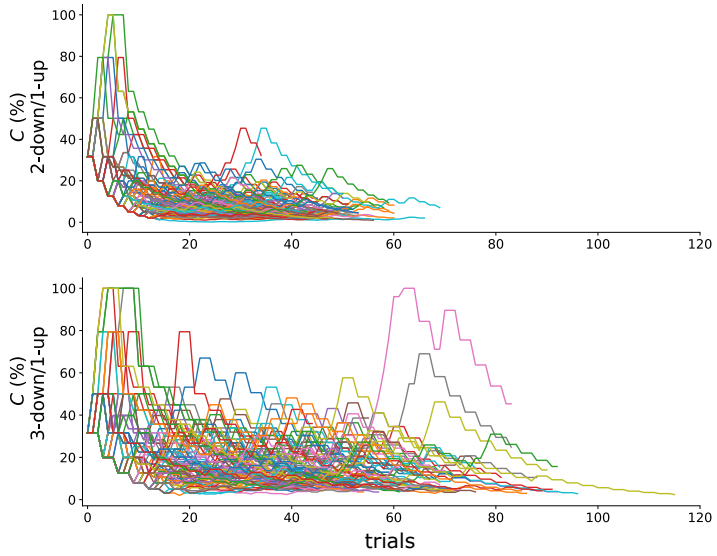

Figure B.1: *Steps of the staircase procedure for the two staircase rules (top: two-up/one-down rule; bottom: three-up/one-down rule). Each line represents one participant.*

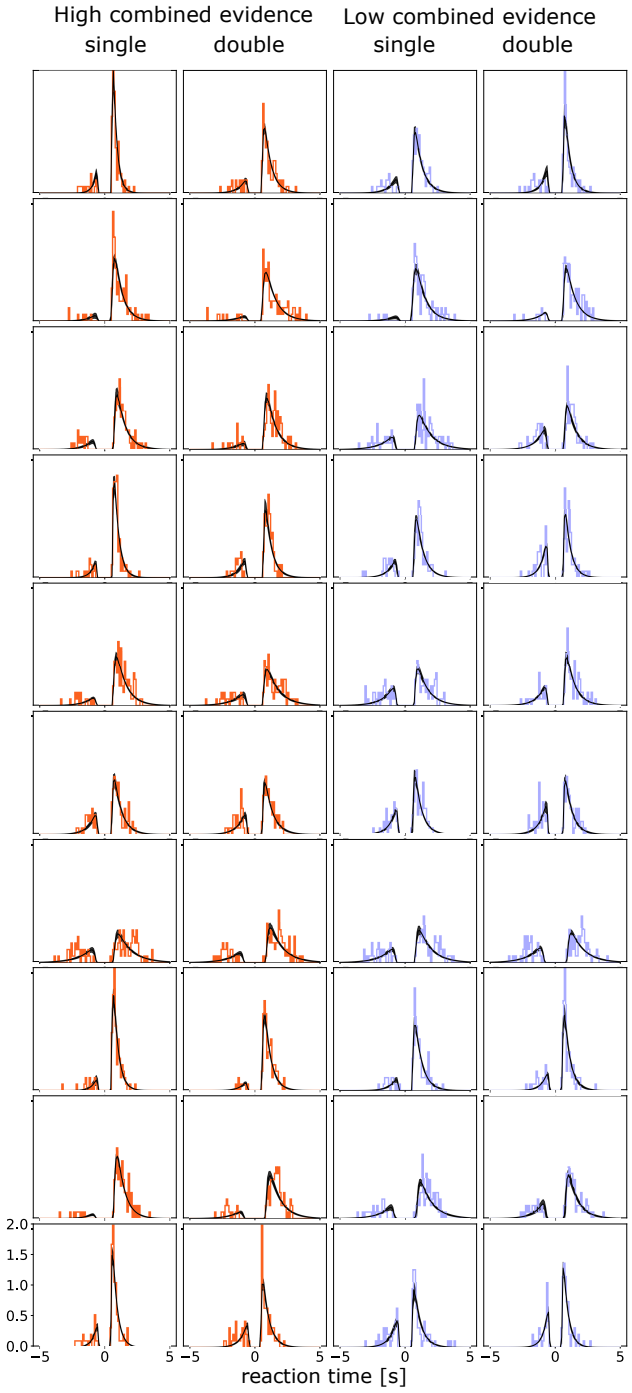

Figure B.2: *Posterior predictive data distributions of 10 participants in Group 1 ($\theta = \pm 20°$). Each row shows data distributions (histograms) as well as Posterior model predictions (black lines) from the best fitted model from one of ten representative participants. The distributions along the positive x-axis indicate normalised correct response times, and the distributions along the negative x-axis indicate normalised error response times. Model predictions was generated in the same procedure as in Figure 4.7.*

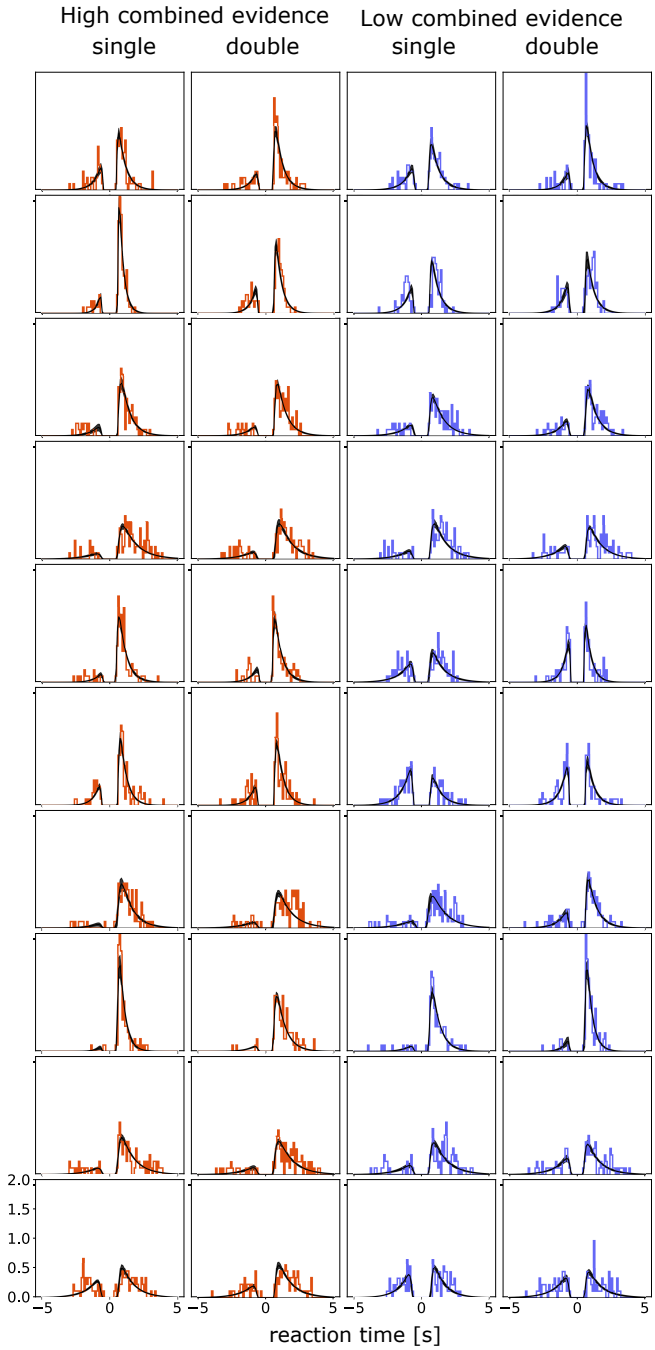

Figure B.3: *Posterior predictive data distributions of 10 participants in Group 2 ($\theta = \pm 45°$). Each row shows data distributions (histograms) as well as Posterior model predictions (black lines) from the best fitted model from one of ten representative participants. The distributions along the positive x-axis indicate normalised correct response times, and the distributions along the negative x-axis indicate normalised error response times. Model predictions was generated in the same procedure as in Figure 4.7*

Bibliography

[1] Birgit Abler, Bärbel Herrnberger, Georg Grön, and Manfred Spitzer. From uncertainty to reward: Bold characteristics differentiate signaling pathways. *BMC neuroscience*, 10 (1):154, 2009.

[2] Kivilcim Afacan-Seref, Natalie A Steinemann, Annabelle Blangero, and Simon P Kelly. Dynamic interplay of value and sensory information in high-speed decision making. *Current Biology*, 28(5):795–802, 2018.

[3] Minkyu Ahn, Sangtae Ahn, Jun Hee Hong, Hohyun Cho, Kiwoong Kim, Bong Soo Kim, Jin Woo Chang, and Sung Chan Jun. Gamma band activity associated with BCI performance: simultaneous MEG/EEG study. *Frontiers in human neuroscience*, 7:848, 2013.

[4] Woo-Young Ahn, Nathaniel Haines, and Lei Zhang. Revealing neurocomputational mechanisms of reinforcement learning and decision-making with the hbayesdm package. *Computational Psychiatry*, 1:24–57, 2017.

[5] V. Aliberti, R. A. Grünewald, C. P. Panayiotopoulos, and E. Chroni. Focal electroencephalographic abnormalities in juvenile myoclonic epilepsy. *Epilepsia*, 35(2):297–301, 1994. doi: 10.1111/j.1528-1157.1994.tb02433.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1528-1157.1994.tb02433.x.

[6] John R Anderson. Act: A simple theory of complex cognition. *American psychologist*, 51(4):355, 1996.

[7] John R Anderson, Michael Matessa, and Christian Lebiere. Act-r: A theory of higher level cognition and its relation to visual attention. *Human–Computer Interaction*, 12 (4):439–462, 1997.

[8] John R Anderson, Christian Lebiere, Marsha Lovett, and Lynne Reder. Act-r: A higher-level account of processing capacity. *Behavioral and Brain Sciences*, 21(6):831–832, 1998.

[9] Philip W Anderson. More is different. *Science*, 177(4047):393–396, 1972.

[10] Jeffrey Annis, Brent J Miller, and Thomas J Palmeri. Bayesian inference with stan: A tutorial on adding custom distributions. *Behavior research methods*, 49(3):863–886, 2017.

[11] Alexander Anwyl-Irvine, Edwin S Dalmaijer, Nick Hodges, and Jo K Evershed. Realistic precision and accuracy of online experiment platforms, web browsers, and devices. *Behavior Research Methods*, pages 1–19, 2020.

[12] Adam R Aron, Trevor W Robbins, and Russell A Poldrack. Inhibition and the right inferior frontal cortex: one decade on. *Trends in cognitive sciences*, 18(4), 2014.

[13] Arian Ashourvan, Shi Gu, Marcelo G. Mattar, Jean M. Vettel, and Danielle S. Bassett. The energy landscape underpinning module dynamics in the human brain connectome. *NeuroImage*, 157:364 – 380, 2017.

[14] Gary Aston-Jones and Jonathan D Cohen. An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.*, 28, 2005.

[15] Nadim AA Atiya, Iñaki Rañó, Girijesh Prasad, and KongFatt Wong-Lin. A neural circuit model of decision uncertainty and change-of-mind. *Nature communications*, 10 (1):1–12, 2019.

[16] R Harald Baayen, Douglas J Davidson, and Douglas M Bates. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of memory and language*, 59(4):390–412, 2008.

[17] A Babloyantz and A Destexhe. Low-dimensional chaos in an instance of epilepsy. *Proceedings of the National Academy of Sciences*, 83(10):3513–3517, 1986. ISSN 0027-8424. doi: 10.1073/pnas.83.10.3513. URL https://www.pnas.org/content/83/10/3513.

[18] R.A.B. Badawy, J. M. Curatolo, M. Newton, S. F. Berkovic, and R. A.L. Macdonell. Sleep deprivation increases cortical excitability in epilepsy. *Neurology*, 67(6):1018–1022, 2006. ISSN 0028-3878. doi: 10.1212/01.wnl.0000237392.64230.f7. URL https://n.neurology.org/content/67/6/1018.

[19] Emilia Bagiella, Richard P Sloan, and Daniel F Heitjan. Mixed-effects models in psychophysiology. *Psychophysiology*, 37(1):13–20, 2000.

[20] Adam P Baker, Matthew J Brookes, Iead A Rezek, Stephen M Smith, Timothy Behrens, Penny J Probert Smith, and Mark Woolrich. Fast transient networks in spontaneous human brain activity. *eLife*, 3:e01867, mar 2014. ISSN 2050-084X. doi: 10.7554/eLife.01867. URL https://doi.org/10.7554/eLife.01867.

[21] Akram Bakkour, Ariel Zylberberg, Michael N Shadlen, and Daphna Shohamy. Value-based decisions involve sequential sampling from memory. *BioRxiv*, page 269290, 2018.

[22] Ian C Ballard, Bokyung Kim, Anthony Liatsis, Gökhan Aydogan, Jonathan D Cohen, and Samuel M McClure. More is meaningful: the magnitude effect in intertemporal choice depends on self-control. *Psychological science*, 28(10):1443–1454, 2017.

[23] Cornelia Bargmann, William Newsome, A Anderson, E Brown, Karl Deisseroth, J Donoghue, Peter MacLeish, E Marder, R Normann, J Sanes, et al. Brain 2025: a scientific vision. *Brain Research through Advancing Innovative Neurotechnologies (BRAIN) Working Group Report to the Advisory Committee to the Director, NIH*, 2014.

[24] Dale J Barr, Roger Levy, Christoph Scheepers, and Harry J Tily. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language*, 68(3):255–278, 2013.

[25] Danielle S Bassett, Perry Zurn, and Joshua I Gold. On the nature and use of models in network neuroscience. *Nature Reviews Neuroscience*, 19(9):566–578, 2018.

[26] Douglas Bates, Reinhold Kliegl, Shravan Vasishth, and Harald Baayen. Parsimonious mixed models. *arXiv preprint arXiv:1506.04967*, 2015.

[27] M Jésus Bayarri and James O Berger. The interplay of bayesian and frequentist analysis. *Statistical Science*, pages 58–80, 2004.

[28] Betül Baykan and Peter Wolf. Juvenile myoclonic epilepsy as a spectrum disorder: A focused review. *Seizure*, 49:36 – 41, 2017. ISSN 1059-1311. doi: https://doi.org/10.1016/j.seizure.2017.05.011. URL http://www.sciencedirect.com/science/article/pii/S1059131117302844.

[29] William Bechtel and Robert C Richardson. *Discovering complexity: Decomposition and localization as strategies in scientific research*. MIT press, 2010.

[30] Oren M. Becker and Martin Karplus. The topology of multidimensional potential energy surfaces: Theory and application to peptide structure and kinetics. *The Journal of Chemical Physics*, 106(4):1495–1517, 1997. doi: 10.1063/1.473299. URL https://doi.org/10.1063/1.473299.

[31] Bianca Beersma, John R Hollenbeck, Stephen E Humphrey, Henry Moon, Donald E Conlon, and Daniel R Ilgen. Cooperation, competition, and team performance: Toward a contingency approach. *Academy of Management Journal*, 46(5):572–590, 2003.

[32] Trevor Bekolay, James Bergstra, Eric Hunsberger, Travis DeWolf, Terrence Stewart, Daniel Rasmussen, Xuan Choo, Aaron Voelker, and Chris Eliasmith. Nengo: a Python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics*, 7 (48):1–13, 2014. ISSN 1662-5196. doi: 10.3389/fninf.2013.00048.

[33] Trevor Bekolay, Mark Laubach, and Chris Eliasmith. A spiking neural integrator model of the adaptive control of action by the medial prefrontal cortex. *The Journal of Neuroscience*, 34(5), 2014.

[34] Anthony J Bell and Terrence J Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159, 1995.

[35] Hans Berger. Über das elektrenkephalogramm des menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, pages 527–570, 1929.

[36] Samuel F. Berkovic, R. Anne Howell, David A. Hay, and John L. Hopper. Epilepsies in twins: Genetics of the major epilepsy syndromes. *Annals of Neurology*, 43(4):435–445, 1998.

[37] Marc G Berman, John Jonides, and Derek Evan Nee. Studying mind and brain with fMRI. *Social cognitive and affective neuroscience*, 1(2):158–161, 2006.

[38] Michael Betancourt. A conceptual introduction to hamiltonian monte carlo. *arXiv preprint arXiv:1701.02434*, 2017.

[39] Luiz Eduardo Betting, Susana Barreto Mory, Li Min Li, Iscia Lopes-Cendes, Marilisa M. Guerreiro, Carlos A.M. Guerreiro, and Fernando Cendes. Voxel-based morphometry in patients with idiopathic generalized epilepsies. *NeuroImage*, 32(2):498 – 502, 2006. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2006.04.174. URL http://www.sciencedirect.com/science/article/pii/S1053811906004460.

[40] William Bialek. Perspectives on theory at the interface of physics and biology. *Reports on Progress in Physics*, 81(1):012601, dec 2017. doi: 10.1088/1361-6633/aa995b. URL https://doi.org/10.1088%2F1361-6633%2Faa995b.

[41] Narcisse P Bichot, Kirk G Thompson, S Chenchal Rao, and Jeffrey D Schall. Reliability of macaque frontal eye field neurons signaling saccade targets during visual search. *Journal of Neuroscience*, 21(2):713–725, 2001.

[42] Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.

[43] Gunnar Blohm, Konrad P. Kording, and Paul R. Schrater. A how-to-model guide for neuroscience. *eNeuro*, 7(1), 2020. doi: 10.1523/ENEURO.0352-19.2019. URL https://www.eneuro.org/content/7/1/ENEURO.0352-19.2019.

[44] Bruce Bobier, Terrence C. Stewart, and Chris Eliasmith. A unifying mechanistic model of selective attention in spiking neurons. *PLOS Computational Biology*, 10(6):1–16, 06 2014. doi: 10.1371/journal.pcbi.1003577. URL https://doi.org/10.1371/journal.pcbi.1003577.

[45] Róbert Bódizs, Ferenc Gombos, and Ilona Kovács. Sleep EEG fingerprints reveal accelerated thalamocortical oscillatory dynamics in williams syndrome. *Research in developmental disabilities*, 33(1):153–164, 2012.

[46] Udo Boehm, Guy E Hawkins, Scott Brown, Hedderik van Rijn, and Eric-Jan Wagenmakers. Of monkeys and men: Impatience in perceptual decision-making. *Psychonomic bulletin & review*, 23(3):738–749, 2016.

[47] Rafal Bogacz. Optimal decision-making theories: linking neurobiology with behaviour. *Trends in cognitive sciences*, 11(3):118–125, 2007.

[48] Rafal Bogacz and Kevin Gurney. The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural computation*, 19(2), 2007.

[49] Rafal Bogacz, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D Cohen. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological review*, 113(4):700, 2006.

[50] Rafal Bogacz, Marius Usher, Jiaxiang Zhang, and James L McClelland. Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1485):1655–1670, 2007.

[51] Jeffrey S Bowers. On the biological plausibility of grandmother cells: Implications for neural network theories in psychology and neuroscience. *Psychological review*, 116(1): 220, 2009.

[52] Rachael Briggs. Normative theories of rational choice: Expected utility. In Edward Zalta, editor, *The Stanford Encyclopedia of Philosophy*. 2017.

[53] Kenneth H. Britten, Michael N. Shadlen, William T. Newsome, and J. Anthony Movshon. Responses of neurons in macaque mt to stochastic motion signals. *Visual Neuroscience*, 10(6):1157–1169, 1993. doi: 10.1017/S0952523800010269.

[54] KH Britten, MN Shadlen, WT Newsome, and JA Movshon. The analysis of visual motion: A comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12(12):4745–4765, 1992. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.12-12-04745.1992. URL https://www.jneurosci.org/content/12/12/4745.

[55] D.E. Broadbent. *Perception and Communication*. Elsevier Science, 2013. ISBN 9781483225821. URL https://books.google.co.uk/books?id=ZCOLBQAAQBAJ.

[56] Verena Brodbeck, Alena Kuhn, Frederic von Wegner, Astrid Morzelewski, Enzo Tagliazucchi, Sergey Borisov, Christoph M. Michel, and Helmut Laufs. EEG microstates of wakefulness and nrem sleep. *NeuroImage*, 62(3):2129 – 2139, 2012. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2012.05.060. URL http://www.sciencedirect.com/science/article/pii/S1053811912005484.

[57] Matthew J. Brookes, Mark Woolrich, Henry Luckhoo, Darren Price, Joanne R. Hale, Mary C. Stephenson, Gareth R. Barnes, Stephen M. Smith, and Peter G. Morris. Investigating the electrophysiological basis of resting state networks using magnetoencephalography. *Proceedings of the National Academy of Sciences*, 108(40):16783–16788, 2011. ISSN 0027-8424. doi: 10.1073/pnas.1112685108. URL https://www.pnas.org/content/108/40/16783.

[58] Scott D Brown and Andrew Heathcote. The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive psychology*, 57(3):153–178, 2008.

[59] Nicolas Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of computational neuroscience*, 8(3):183–208, 2000.

[60] Stephen G Brush. History of the Lenz-Ising model. *Reviews of Modern Physics*, 39(4): 883, 1967.

[61] E. Bullmore and O. Sporns. Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci*, 10(3):186–198, Mar 2009.

[62] Jerome R. Busemeyer, Sebastian Gluth, Jörg Rieskamp, and Brandon M. Turner. Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends in Cognitive Sciences*, 23(3):251–263, Mar 2019. ISSN 1364-6613. doi: 10.1016/j.tics.2018.12.003. URL https://doi.org/10.1016/j.tics.2018.12.003.

[63] K. Caeyenberghs, H.W.R. Powell, R.H. Thomas, L. Brindley, C. Church, J. Evans, S.D. Muthukumaraswamy, D.K. Jones, and K. Hamandi. Hyperconnectivity in juvenile myoclonic epilepsy: A network analysis. *NeuroImage: Clinical*, 7:98 – 104, 2015. ISSN 2213-1582. doi: https://doi.org/10.1016/j.nicl.2014.11.018. URL http://www.sciencedirect.com/science/article/pii/S221315821400182X.

[64] Carol S Camfield, Pasquale Striano, and Peter R Camfield. Epidemiology of juvenile myoclonic epilepsy. *Epilepsy & Behavior*, 28:S15–S17, 2013.

[65] Nicholas T Carnevale and Michael L Hines. *The NEURON book*. Cambridge Univ. Press, 2006.

[66] Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of statistical software*, 76(1), 2017.

[67] James F Cavanagh, Thomas V Wiecki, Michael X Cohen, Christina M Figueroa, Johan Samanta, Scott J Sherman, and Michael J Frank. Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature neuroscience*, 14(11): 1462, 2011.

[68] Samuel R Chamberlain, Adam Hampshire, Ulrich Müller, Katya Rubia, Natalia Del Campo, Kevin Craig, Ralf Regenthal, John Suckling, Jonathan P Roiser, Jon E Grant, et al. Atomoxetine modulates right inferior frontal activation during inhibitory control: a pharmacological functional magnetic resonance imaging study. *Biological Psychiatry*, 65(7), 2009.

[69] Lucie Charles, Camille Chardin, and Patrick Haggard. Evidence for metacognitive bias in perception of voluntary action. *Cognition*, 194:104041, 2020.

[70] Xing-jie Chen and Youngbin Kwak. What makes you go faster?: the effect of reward on speeded action under risk. *Frontiers in psychology*, 8:1057, 2017.

[71] Feng-Xuan Choo and Chris Eliasmith. A spiking neuron model of serial-order recall. In *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*, 2010.

[72] Fahmida A. Chowdhury, Wessel Woldman, Thomas H. B. FitzGerald, Robert D. C. Elwes, Lina Nashef, John R. Terry, and Mark P. Richardson. Revealing a brain network endophenotype in families with idiopathic generalised epilepsy. *PLOS ONE*, 9(10): 1–8, 10 2014. doi: 10.1371/journal.pone.0110136. URL https://doi.org/10.1371/journal.pone.0110136.

[73] Radoslaw Martin Cichy, Dimitrios Pantazis, and Aude Oliva. Resolving human object recognition in space and time. *Nature neuroscience*, 17(3):455, 2014.

[74] Emilio N. M. Cirillo and Joel L. Lebowitz. Metastability in the two-dimensional ising model with free boundary conditions. *Journal of Statistical Physics*, 90(1):211–226, Jan 1998. ISSN 1572-9613. doi: 10.1023/A:1023255802455. URL https://doi.org/10.1023/A:1023255802455.

[75] Paul Cisek. Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362 (1485):1585–1599, 2007.

[76] B. Clemens, S. Puskás, M. Besenyei, T. Spisák, G. Opposits, K. Hollódy, A. Fogarasi, I. Fekete, and M. Emri. Neurophysiology of juvenile myoclonic epilepsy: EEG-based network and graph analysis of the interictal and immediate preictal states. *Epilepsy Research*, 106(3):357 – 369, 2013. ISSN 0920-1211. doi: https://doi.org/10.1016/j.eplepsyres.2013.06.017. URL http://www.sciencedirect.com/science/article/pii/S0920121113001812.

[77] Béla Clemens. Pathological theta oscillations in idiopathic generalised epilepsy. *Clinical Neurophysiology*, 115(6):1436 – 1441, 2004. ISSN 1388-2457. doi: https://doi.org/10.1016/j.clinph.2004.01.018. URL http://www.sciencedirect.com/science/article/pii/S1388245704000379.

[78] Anton Coenen, Edward Fine, and Oksana Zayachkivska. Adolf Beck: A forgotten pioneer in electroencephalography. *Journal of the History of the Neurosciences*, 23(3): 276–286, 2014. doi: 10.1080/0964704X.2013.867600. URL https://doi.org/10.1080/0964704X.2013.867600. PMID: 24735457.

[79] Andrew L Cohen, Namyi Kang, and Tanya L Leise. Multi-attribute, multi-alternative models of choice: Choice, reaction time, and process tracing. *Cognitive psychology*, 98: 45–72, 2017.

[80] Steven J Cooper. Donald O. Hebb's synapse and learning rule: a history and commentary. *Neuroscience & Biobehavioral Reviews*, 28(8):851–874, 2005.

[81] F. H. L. da Silva, W. Blanes, S. N. Kalitzin, J. Parra, P. Suffczynski, and D. N. Velis. Dynamical diseases of brain systems: different routes to epileptic seizures. *IEEE Transactions on Biomedical Engineering*, 50(5):540–548, May 2003. ISSN 0018-9294. doi: 10.1109/TBME.2003.810703.

[82] Avijit Datta, Rhodri Cusack, Kari Hawkins, Joost Heutink, Chris Rorden, Ian H Robertson, and Tom Manly. The p300 as a marker of waning attention and error propensity. *Computational intelligence and neuroscience*, 2007, 2007.

[83] Olivier David, James M Kilner, and Karl J Friston. Mechanisms of evoked and induced responses in MEG/EEG. *Neuroimage*, 31(4):1580–1591, 2006.

[84] Peter Dayan and Laurence F Abbott. *Theoretical neuroscience: computational and mathematical modeling of neural systems*. MIT Press, 2001.

[85] Joshua R. de Leeuw. jsPsych: A javascript library for creating behavioral experiments in a web browser. *Behavior Research Methods*, 47(1):1–12, Mar 2015. ISSN 1554-3528. doi: 10.3758/s13428-014-0458-y. URL https://doi.org/10.3758/s13428-014-0458-y.

[86] Joshua R. de Leeuw and Benjamin A. Motz. Psychophysics in a web browser? Comparing response times collected with javascript and psychophysics toolbox in a visual search task. *Behavior Research Methods*, 48(1):1–12, Mar 2016. ISSN 1554-3528. doi: 10.3758/s13428-015-0567-2. URL https://doi.org/10.3758/s13428-015-0567-2.

[87] Gustavo Deco, Viktor K. Jirsa, Peter A. Robinson, Michael Breakspear, and Karl Friston. The dynamic brain: From spiking neurons to neural masses and cortical fields. *PLOS Computational Biology*, 4(8):1–35, 08 2008. doi: 10.1371/journal.pcbi.1000092. URL https://doi.org/10.1371/journal.pcbi.1000092.

[88] Gustavo Deco, Mario Senden, and Viktor Jirsa. How anatomy shapes dynamics: a semi-analytical study of the brain at rest by a simple spin model. *Frontiers in Computational Neuroscience*, 6:68, 2012. ISSN 1662-5188. doi: 10.3389/fncom.2012.00068. URL https://www.frontiersin.org/article/10.3389/fncom.2012.00068.

[89] Gustavo Deco, Adrián Ponce-Alvarez, Dante Mantini, Gian Luca Romani, Patric Hagmann, and Maurizio Corbetta. Resting-state functional connectivity emerges from structurally and dynamically shaped slow linear fluctuations. *Journal of Neuroscience*, 33 (27):11239–11252, 2013.

[90] A. V. Delgado-Escueta and Fe Enrile-Bacsal. Juvenile myoclonic epilepsy of janz. *Neurology*, 34(3):285–285, 1984. ISSN 0028-3878. doi: 10.1212/WNL.34.3.285. URL http://n.neurology.org/content/34/3/285.

[91] Arnaud Delorme and Scott Makeig. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of neuroscience methods*, 134(1):9–21, 2004.

[92] Kobe Desender, Annika Boldt, Tom Verguts, and Tobias H Donner. Confidence predicts speed-accuracy tradeoff for subsequent decisions. *Elife*, 8:e43499, 2019.

[93] Rahul S. Desikan, Florent Ségonne, Bruce Fischl, Brian T. Quinn, Bradford C. Dickerson, Deborah Blacker, Randy L. Buckner, Anders M. Dale, R. Paul Maguire, Bradley T. Hyman, Marilyn S. Albert, and Ronald J. Killiany. An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. *NeuroImage*, 31(3):968 – 980, 2006. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2006.01.021. URL http://www.sciencedirect.com/science/article/pii/S1053811906000437.

[94] Mark D'Esposito. From cognitive to neural models of working memory. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 2007.

[95] Christophe Destrieux, Bruce Fischl, Anders Dale, and Eric Halgren. Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. *NeuroImage*, 53(1):1 – 15, 2010. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2010.06.010. URL http://www.sciencedirect.com/science/article/pii/S1053811910008542.

[96] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[97] Diana C Dima, Gavin Perry, Eirini Messaritaki, Jiaxiang Zhang, and Krish D Singh. Spatiotemporal dynamics in human visual cortex rapidly encode the emotional content of faces. *Human brain mapping*, 39(10):3993–4006, 2018.

[98] Jochen Ditterich. Evidence for time-variant decision making. *European Journal of Neuroscience*, 24(12):3628–3641, 2006.

[99] Jochen Ditterich, Mark E Mazurek, and Michael N Shadlen. Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature neuroscience*, 6(8):891–898, 2003.

[100] Chris Donkin, Scott Brown, Andrew Heathcote, and Eric-Jan Wagenmakers. Diffusion versus linear ballistic accumulation: Different models but the same conclusions about

psychological processes? *Psychonomic Bulletin & Review*, 18(1):61–69, 2011. doi: 10.3758/s13423-010-0022-4. URL https://doi.org/10.3758/s13423-010-0022-4.

[101] Pamela K Douglas and Ariana Anderson. *Feature Fallacy: Complications with Interpreting Linear Decoding Weights in fMRI*, pages 363–378. Springer-Verlang, Berlin, 2019. ISBN 978-3-030-28954-6. doi: 10.1007/978-3-030-28954-6_20. URL https://doi.org/10.1007/978-3-030-28954-6_20.

[102] Pamela K Douglas and Farzad Vasheghani Farahani. On the similarity of deep learning representations across didactic and adversarial examples. *arXiv preprint arXiv:2002.06816*, 2020.

[103] Peter Duggins, D. Krzeminski, C. Eliasmith, and S. Wichary. A spiking neuron model of inferential decision making: Urgency, uncertainty, and the speed-accuracy tradeoff. 42nd Cognitive Science Annual Conference, Toronto, 2020. Cognitive Science Society, 2020.

[104] Gilles Dutilh and Jörg Rieskamp. Comparing perceptual and preferential decision making. *Psychonomic Bulletin & Review*, 23(3):723–737, Jun 2016. ISSN 1531-5320. doi: 10.3758/s13423-015-0941-1. URL https://doi.org/10.3758/s13423-015-0941-1.

[105] Gilles Dutilh, Jeffrey Annis, Scott D. Brown, . . . , Thomas V. Wiecki, Jörg Rieskamp, and Chris Donkin. The quality of response time data inference: A blinded, collaborative assessment of the validity of cognitive models. *Psychonomic Bulletin & Review*, 26 (4):1051–1069, Aug 2019. ISSN 1531-5320. doi: 10.3758/s13423-017-1417-2. URL https://doi.org/10.3758/s13423-017-1417-2.

[106] Philip Eckhoff, KongFatt Wong-Lin, and Philip Holmes. Optimality and robustness of a biophysical decision-making model under norepinephrine modulation. *Journal of Neuroscience*, 29(13):4301–4311, 2009.

[107] Philip Eckhoff, KongFatt Wong-Lin, and Philip Holmes. Dimension reduction and dynamics of a spiking neural network model for decision making under neuromodulation. *SIAM journal on applied dynamical systems*, 10(1):148–188, 2011.

[108] Martin Eimer. The n2pc component as an indicator of attentional selectivity. *Electroencephalography and clinical neurophysiology*, 99(3):225–234, 1996.

[109] Gaute T Einevoll, Alain Destexhe, Markus Diesmann, Sonja Grün, Viktor Jirsa, Marc de Kamps, Michele Migliore, Torbjørn V Ness, Hans E Plesser, and Felix Schürmann. The scientific case for brain simulations. *Neuron*, 102(4):735–744, 2019.

[110] Anders Eklund, Thomas E Nichols, and Hans Knutsson. Cluster failure: Why fmri inferences for spatial extent have inflated false-positive rates. *Proceedings of the national academy of sciences*, 113(28):7900–7905, 2016.

[111] C. E. Elger, G. Widman, R. Andrzejak, J. Arnhold, P. David, and K. Lehnertz. Nonlinear EEG analysis and its potential role in epileptology. *Epilepsia*, 41(s3):S34–S38, 2000. doi: 10.1111/j.1528-1157.2000.tb01532.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1528-1157.2000.tb01532.x.

[112] Chris Eliasmith. The third contender: A critical examination of the dynamicist theory of cognition. *Philosophical Psychology*, 9(4):441–463, 1996.

[113] Chris Eliasmith. *How to build a brain: A neural architecture for biological cognition*. Oxford Univ. Press, 2013.

[114] Chris Eliasmith and Charles H Anderson. *Neural engineering: Computation, Representation, and Dynamics in Neurobiological Systems*. MIT press, 2003.

[115] Adham Elshahabi, Silke Klamer, Ashish Kaul Sahib, Holger Lerche, Christoph Braun, and Niels K. Focke. Magnetoencephalography reveals a widespread increase in network connectivity in idiopathic/genetic generalized epilepsy. *PLOS ONE*, 10(9):1–16, 09 2015. doi: 10.1371/journal.pone.0138119. URL https://doi.org/10.1371/journal.pone.0138119.

[116] Pedley TA Engel J. *Epilepsy: A comprehensive textbook*. Lippincott-Raven, 1997.

[117] G Bard Ermentrout and Nancy Kopell. Parabolic bursting in an excitable system coupled with a slow oscillation. *SIAM Journal on Applied Mathematics*, 46(2):233–253, 1986.

[118] Guillem R Esber and Mark Haselgrove. Reconciling the influence of predictiveness and uncertainty on stimulus salience: a model of attention in associative learning. *Proceedings of the Royal Society B: Biological Sciences*, 278(1718):2553–2561, 2011.

[119] Nathan Evans. Think fast! the implications of emphasizing urgency in decision-making. *psyarchiv*, 2020. doi: 10.31234/osf.io/pfrb4.

[120] Takahiro Ezaki, Michiko Sakaki, Takamitsu Watanabe, and Naoki Masuda. Age-related changes in the ease of dynamical transitions in human brain activity. *Human Brain Mapping*, 39(6):2673–2688, 2018. doi: 10.1002/hbm.24033. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.24033.

[121] Takahiro Ezaki, Elohim Fonseca Dos Reis, Takamitsu Watanabe, Michiko Sakaki, and Naoki Masuda. Closer to critical resting-state neural dynamics in individuals with higher fluid intelligence. *Communications biology*, 3(1):1–9, 2020.

[122] Daniele Fanelli. Opinion: Is science really facing a reproducibility crisis, and do we need it to? *Proceedings of the National Academy of Sciences*, 115(11):2628–2631, 2018.

[123] Santiago Fernandez, Antonio Donaire, Iratxe Maestro, Eulalia Seres, Xavier Setoain, Nuria Bargalló, Jordi Rumià, Teresa Boget, Carles Falcón, and Mar Carreño. Functional neuroimaging in startle epilepsy: Involvement of a mesial frontoparietal network. *Epilepsia*, 52(9):1725–1732, 2011. doi: 10.1111/j.1528-1167.2011.03172.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1528-1167.2011.03172.x.

[124] Christopher D Fiorillo, Philippe N Tobler, and Wolfram Schultz. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, 299(5614):1898–1902, 2003.

[125] Peter C Fishburn. Normative theories of decision making under risk and under uncertainty. In *Non-Conventional Preference Relations in Decision Making*, pages 1–21. Springer, 1988.

[126] Richard FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical journal*, 1(6):445, 1961.

[127] Jerry A. Fodor. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press, 1987.

[128] Birte U Forstmann, Gilles Dutilh, Scott Brown, Jane Neumann, D Yves Von Cramon, K Richard Ridderinkhof, and Eric-Jan Wagenmakers. Striatum and pre-SMA facilitate

decision-making under time pressure. *Proceedings of the National Academy of Sciences*, 105(45):17538–17542, 2008.

[129] Birte U Forstmann, Alfred Anwander, Andreas Schäfer, Jane Neumann, Scott Brown, Eric-Jan Wagenmakers, Rafal Bogacz, and Robert Turner. Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences*, 107(36), 2010.

[130] Michael J Frank, Chris Gagne, Erika Nyhus, Sean Masters, Thomas Wiecki, James F Cavanagh, and David Badre. fMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *Journal of Neuroscience*, 35(2):485–494, 2015.

[131] Yves Frégnac and Gilles Laurent. Neuroscience: Where is the brain in the human brain project? *Nature News*, 513(7516):27, 2014.

[132] Karl J Friston, Lee Harrison, and Will Penny. Dynamic causal modelling. *Neuroimage*, 19(4):1273–1302, 2003.

[133] Miguel A Garcıa-Pérez. Forced-choice staircases with fixed step sizes: asymptotic and small-sample properties. *Vision research*, 38(12):1861–1881, 1998.

[134] Miguel A. García-Pérez. Optimal setups for forced-choice staircases with fixed step sizes. *Spatial Vision*, 13(4):431 – 448, 2000. doi: https://doi.org/10.1163/156856800741306. URL https://brill.com/view/journals/sv/13/4/article-p431_7.xml.

[135] Deon Garrett, David A Peterson, Charles W Anderson, and Michael H Thaut. Comparison of linear, nonlinear, and feature selection methods for EEG signal classification. *IEEE Transactions on neural systems and rehabilitation engineering*, 11(2):141–144, 2003.

[136] M. Gazzaniga and R.B. Ivry. *Cognitive Neuroscience: The Biology of the Mind: Fourth International Student Edition*. International student edition. Norton, 2013. ISBN 9780393922288. URL https://books.google.co.uk/books?id=bPvdCgAAQBAJ.

[137] Andrew Gelman and Donald B Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7:457–472, 1992. URL http://www.jstor.org/about/terms.html.

[138] Andrew Gelman, Donald B Rubin, et al. Inference from iterative simulation using multiple sequences. *Statistical science*, 7(4):457–472, 1992.

[139] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian Data Analysis*. CRC press, 2013.

[140] Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. *Neuronal dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge Univ. Press, 2014.

[141] Gerd Gigerenzer and Peter M Todd. Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart*. Oxford Univ. Press, 1999.

[142] Joshua I Glaser, Ari S Benjamin, Roozbeh Farhoodi, and Konrad P Kording. The roles of supervised machine learning in systems neuroscience. *Progress in Neurobiology*, 175: 126–137, 2019.

[143] Joshua I Glaser, Ari S Benjamin, Raeed H Chowdhury, Matthew G Perich, Lee E Miller, and Konrad P Kording. Machine learning for neural decoding. *Eneuro*, 7(4), 2020.

[144] Nigel H Goddard, Michael Hucka, Fred Howell, Hugo Cornelis, Kavita Shankar, and David Beeman. Towards neuroml: model description methods for collaborative modelling in neuroscience. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 356(1412):1209–1228, 2001.

[145] Imran Goker, Onur Osman, Serhat Ozekes, M. Baris Baslo, Mustafa Ertas, and Yekta Ulgen. Classification of juvenile myoclonic epilepsy data acquired through scanning electromyography with machine learning algorithms. *Journal of Medical Systems*, 36 (5):2705–2711, Oct 2012. ISSN 1573-689X. doi: 10.1007/s10916-011-9746-6. URL https://doi.org/10.1007/s10916-011-9746-6.

[146] Mateusz Gola, Mikołaj Magnuski, Izabela Szumska, and Andrzej Wróbel. EEG beta band activity is related to attention and attentional deficits in the visual performance of elderly subjects. *International Journal of Psychophysiology*, 89(3):334–341, 2013.

[147] Joshua I Gold and Michael N Shadlen. Neural computations that underlie decisions about sensory stimuli. *Trends in cognitive sciences*, 5(1):10–16, 2001.

[148] Joshua I Gold and Michael N Shadlen. The neural basis of decision making. *Annual review of neuroscience*, 30, 2007.

[149] Stephanie Goldfarb, Naomi E. Leonard, Patrick Simen, Carlos H. Caicedo-Núñez, and Philip Holmes. A comparative study of drift diffusion and linear ballistic accumulator models in a reward maximization perceptual choice task. *Frontiers in neuroscience*, 8:148–148, Aug 2014. ISSN 1662-4548. doi: 10.3389/fnins.2014.00148. URL https://pubmed.ncbi.nlm.nih.gov/25140124. 25140124[pmid].

[150] Robin I Goldman, John M Stern, Jerome Engel Jr, and Mark S Cohen. Simultaneous EEG and fMRI of the alpha rhythm. *Neuroreport*, 13(18):2487, 2002.

[151] Evan M. Gordon, Timothy O. Laumann, Babatunde Adeyemo, Jeremy F. Huckins, William M. Kelley, and Steven E. Petersen. Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. *Cerebral Cortex*, 26(1):288–303, 10 2014. ISSN 1047-3211. doi: 10.1093/cercor/bhu239. URL https://doi.org/10.1093/cercor/bhu239.

[152] Krzysztof J Gorgolewski, Fidel Alfaro-Almagro, Tibor Auer, Pierre Bellec, Mihai Capotă, M Mallar Chakravarty, Nathan W Churchill, Alexander Li Cohen, R Cameron Craddock, Gabriel A Devenyi, et al. Bids apps: Improving ease of use, accessibility, and reproducibility of neuroimaging data analysis methods. *PLoS computational biology*, 13 (3):e1005209, 2017.

[153] J. Gotman, C. Grova, A. Bagshaw, E. Kobayashi, Y. Aghakhani, and F. Dubeau. Generalized epileptic discharges show thalamocortical activation and suspension of the default state of the brain. *Proceedings of the National Academy of Sciences*, 102(42):15236–15240, 2005. ISSN 0027-8424. doi: 10.1073/pnas.0504935102. URL https://www.pnas.org/content/102/42/15236.

[154] David Marvin Green, John A Swets, et al. *Signal Detection Theory and Psychophysics*, volume 1. Wiley, New York, 1966.

[155] Shi Gu, Matthew Cieslak, Benjamin Baird, Sarah F. Muldoon, Scott T. Grafton, Fabio Pasqualetti, and Danielle S. Bassett. The energy landscape of neurophysiological activity implicit in brain network structure. *Scientific Reports*, 8(1):2507, 2018.

ISSN 2045-2322. doi: 10.1038/s41598-018-20123-8. URL https://doi.org/10.1038/s41598-018-20123-8.

[156] Ralitza Gueorguieva and John H Krystal. Move over anova: progress in analyzing repeated-measures data andits reflection in papers published in the archives of general psychiatry. *Archives of general psychiatry*, 61(3):310–317, 2004.

[157] Emma L. Hall, Mark W. Woolrich, Carlos E. Thomaz, Peter G. Morris, and Matthew J. Brookes. Using variance information in magnetoencephalography measures of functional connectivity. *NeuroImage*, 67:203 – 212, 2013. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2012.11.011. URL http://www.sciencedirect.com/science/article/pii/S105381191201107X.

[158] Khalid Hamandi, Afraim Salek-Haddadi, Helmut Laufs, Adam Liston, Karl Friston, David R. Fish, John S. Duncan, and Louis Lemieux. EEG–fMRI of idiopathic and secondarily generalized epilepsies. *NeuroImage*, 31(4):1700 – 1710, 2006. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2006.02.016. URL http://www.sciencedirect.com/science/article/pii/S1053811906001108.

[159] Timothy D. Hanks and Christopher Summerfield. Perceptual decision making in rodents, monkeys, and humans. *Neuron*, 93:15–31, 1 2017. ISSN 10974199. doi: 10.1016/j.neuron.2016.12.003.

[160] Timothy D Hanks, Charles D Kopec, Bingni W Brunton, Chunyu A Duan, Jeffrey C Erlich, and Carlos D Brody. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*, 520(7546):220–223, 2015.

[161] Enrique CA Hansen, Demian Battaglia, Andreas Spiegler, Gustavo Deco, and Viktor K Jirsa. Functional connectivity dynamics: modeling the switching behavior of the resting state. *Neuroimage*, 105:525–535, 2015.

[162] Peter Hansen, Morten Kringelbach, and Riitta Salmelin. *MEG: An Introduction to Methods*. Oxford Univ. Press, 2010.

[163] Riitta Hari and Aina Puce. *MEG-EEG Primer*. Oxford Univ. Press, 2017.

[164] Stevan Harnad. The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42 (1-3):335–346, 1990.

[165] Demis Hassabis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. Neuroscience-inspired artificial intelligence. *Neuron*, 95(2):245–258, 2017.

[166] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 04 1970. ISSN 0006-3444. doi: 10.1093/biomet/57.1.97. URL https://doi.org/10.1093/biomet/57.1.97.

[167] Elke Hattingen, Christian Lückerath, Stefanie Pellikan, Dmitri Vronski, Christine Roth, Susanne Knake, Matthias Kieslich, and Ulrich Pilatus. Frontal and thalamic changes of gaba concentration indicate dysfunction of thalamofrontal networks in juvenile myoclonic epilepsy. *Epilepsia*, 55(7):1030–1037, 2014. doi: 10.1111/epi.12656. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/epi.12656.

[168] Stefan Haufe, Frank Meinecke, Kai Görgen, Sven Dähne, John-Dylan Haynes, Benjamin Blankertz, and Felix Bießmann. On the interpretation of weight vectors of linear models in multivariate neuroimaging. *Neuroimage*, 87:96–110, 2014.

[169] Guy E Hawkins, Eric-Jan Wagenmakers, Roger Ratcliff, and Scott D Brown. Discriminating evidence accumulation from urgency signals in speeded decision making, 2015.

[170] Hauke R Heekeren, Sean Marrett, Peter A Bandettini, and Leslie G Ungerleider. A general mechanism for perceptual decision-making in the human brain. *Nature*, 431 (7010):859–862, 2004.

[171] Hauke R Heekeren, Sean Marrett, and Leslie G Ungerleider. The neural systems that mediate human perceptual decision making. *Nature reviews neuroscience*, 9(6):467–479, 2008.

[172] Richard P Heitz. The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Frontiers in neuroscience*, 8:150, 2014.

[173] Richard P Heitz and Jeffrey D Schall. Neural mechanisms of speed-accuracy tradeoff. *Neuron*, 76(3):616–628, 2012.

[174] Moritz Helmstaedter. The mutual inspirations of machine learning and neuroscience. *Neuron*, 86(1):25–28, 2015.

[175] Adrián Hernández, Antonio Zainos, and Ranulfo Romo. Temporal evolution of a decision-making process in medial premotor cortex. *Neuron*, 33(6):959–972, 2002.

[176] William E Hick. On the rate of gain of information. *Quarterly Journal of experimental psychology*, 4(1):11–26, 1952.

[177] S. Hill and H. Markram. The Blue Brain Project. In *30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, page 158, 2008. doi: 10.1109/IEMBS.2008.4649065.

[178] Arjan Hillebrand and Gareth R Barnes. Beamformer analysis of MEG data. *International review of neurobiology*, 68:149–171, 2005.

[179] Arjan Hillebrand, Prejaas Tewarie, Edwin van Dellen, Meichen Yu, Ellen W. S. Carbo, Linda Douw, Alida A. Gouw, Elisabeth C. W. van Straaten, and Cornelis J. Stam. Direction of information flow in large-scale resting-state networks is frequency-dependent. *Proceedings of the National Academy of Sciences*, 113(14):3867–3872, 2016. ISSN 0027-8424. doi: 10.1073/pnas.1515657113. URL https://www.pnas.org/content/113/14/3867.

[180] Joerg F. Hipp, David J. Hawellek, Maurizio Corbetta, Markus Siegel, and Andreas K. Engel. Large-scale cortical correlation structure of spontaneous oscillatory activity. *Nature Neuroscience*, 15:884, May 2012. URL https://doi.org/10.1038/nn.3101. Article.

[181] Alan L Hodgkin and Andrew F Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4):500, 1952.

[182] Alan L Hodgkin, Andrew F Huxley, and Bernard Katz. Measurement of current-voltage relations in the membrane of the giant axon of loligo. *The Journal of physiology*, 116 (4):424, 1952.

[183] Andreas Holzinger, Georg Langs, Helmut Denk, Kurt Zatloukal, and Heimo Müller. Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(4):e1312, 2019.

[184] C. J. Honey, O. Sporns, L. Cammoun, X. Gigandet, J. P. Thiran, R. Meuli, and P. Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040, 2009. ISSN 0027-8424. doi: 10.1073/pnas.0811168106. URL https://www.pnas.org/content/106/6/2035.

[185] Barry Horwitz. The elusive concept of brain connectivity. *NeuroImage*, 19(2):466 – 470, 2003. ISSN 1053-8119. doi: https://doi.org/10.1016/S1053-8119(03)00112-5. URL http://www.sciencedirect.com/science/article/pii/S1053811903001125.

[186] Gregory D. Horwitz and William T. Newsome. Separate signals for target selection and movement specification in the superior colliculus. *Science*, 284(5417):1158–1161, 1999. ISSN 0036-8075. doi: 10.1126/science.284.5417.1158. URL https://science.sciencemag.org/content/284/5417/1158.

[187] Robert S Hurley, Borna Bonakdarpour, Xue Wang, and M-Marsel Mesulam. Asymmetric connectivity between the anterior temporal lobe and the language network. *Journal of cognitive neuroscience*, 27(3):464–473, 2015.

[188] Aapo Hyvarinen. Fast ICA for noisy data using Gaussian moments. In *1999 IEEE International Symposium on Circuits and Systems (ISCAS)*, volume 5, pages 57–61. IEEE, 1999.

[189] Leonidas D. Iasemidis, J. Chris Sackellares, Hitten P. Zaveri, and William J. Williams. Phase space topography and the lyapunov exponent of electrocorticograms in partial seizures. *Brain Topography*, 2(3):187–201, Mar 1990. doi: 10.1007/BF01140588. URL https://doi.org/10.1007/BF01140588.

[190] Satoshi Ikemoto. Brain reward circuitry beyond the mesolimbic dopamine system: A neurobiological theory. *Neuroscience & biobehavioral reviews*, 35(2):129–150, 2010.

[191] Jason Ivanoff, Philip Branning, and René Marois. fmri evidence for a dual process account of the speed-accuracy tradeoff in decision-making. *PLoS one*, 3(7):e2635, 2008.

[192] Eugene M Izhikevich. Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks*, 15(5):1063–1070, 2004.

[193] Eugene M Izhikevich. *Dynamical Systems in Neuroscience*. MIT Press, 2007.

[194] Keise Izuma, Madoka Matsumoto, Kou Murayama, Kazuyuki Samejima, Norihiro Sadato, and Kenji Matsumoto. Neural correlates of cognitive dissonance and choice-induced preference change. *Proceedings of the National Academy of Sciences*, 107(51): 22014–22019, 2010.

[195] Sara Jahfari, K Richard Ridderinkhof, and H Steven Scholte. Spatial frequency information modulates response inhibition and decision-making processes. *PLoS One*, 8(10), 2013.

[196] William James. *The principles of psychology*. New York: Holt, 1890.

[197] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review. Series II*, 106(4):620 – 630, 1957. doi: 10.1103/PhysRev.106.620.

[198] Ke Jia, Xin Xue, Jong-Hwan Lee, Fang Fang, Jiaxiang Zhang, and Sheng Li. Visual perceptual learning modulates decision network in the human brain: The evidence from psychophysics, modeling, and functional magnetic resonance imaging. *Journal of Vision*, 18(12):9–9, 2018.

[199] Viktor K Jirsa, T Proix, Dionysios Perdikis, Michael Marmaduke Woodman, H Wang, J Gonzalez-Martinez, Christophe Bernard, Christian Bénar, Maxime Guye, Patrick Chauvel, et al. The virtual epileptic patient: individualized whole-brain models of epilepsy spread. *Neuroimage*, 145:377–388, 2017.

[200] Marcia K Johnson, Carol L Raye, Karen J Mitchell, Erich J Greene, William A Cunningham, and Charles A Sanislow. Using fMRI to investigate. *Cognitive, Affective, & Behavioral Neuroscience*, 5(3):339–361, 2005.

[201] T-P Jung, Colin Humphries, T-W Lee, Scott Makeig, Martin J McKeown, Vicente Iragui, and Terrence J Sejnowski. Removing electroencephalographic artifacts: comparison between ICA and PCA. In *Neural Networks for Signal Processing VIII. Proceedings of the 1998 IEEE Signal Processing Society Workshop*, pages 63–72. IEEE, 1998.

[202] Leo P. Kadanoff. More is the Same; Phase Transitions and Mean Field Theories. *Journal of Statistical Physics*, 137(5-6):777–797, December 2009. doi: 10.1007/s10955-009-9814-1.

[203] Thorsten Kahnt, Jakob Heinzle, Soyoung Q. Park, and John-Dylan Haynes. Decoding different roles for vmpfc and dlpfc in multi-attribute decision making. *NeuroImage*, 56(2):709 – 715, 2011. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2010.05.058. URL http://www.sciencedirect.com/science/article/pii/S1053811910007913. Multivariate Decoding and Brain Reading.

[204] Jiyoung Kang, Chongwon Pae, and Hae-Jeong Park. Graph-theoretical analysis for energy landscape reveals the organization of state transitions in the resting-state human cerebral cortex. *PLOS ONE*, 14(9):1–25, 09 2019. doi: 10.1371/journal.pone.0222161. URL https://doi.org/10.1371/journal.pone.0222161.

[205] N. Kannathal, Min Lim Choo, U. Rajendra Acharya, and P.K. Sadasivan. Entropies for detection of epilepsy in EEG. *Computer Methods and Programs in Biomedicine*, 80 (3):187 – 194, 2005. ISSN 0169-2607. doi: https://doi.org/10.1016/j.cmpb.2005.06.012. URL http://www.sciencedirect.com/science/article/pii/S0169260705001525.

[206] Esin Karahan, Alison G Costigan, Kim S Graham, Andrew D Lawrence, and Jiaxiang Zhang. Cognitive and white-matter compartment models reveal selective relations between corticospinal tract microstructure and simple reaction time. *Journal of Neuroscience*, 39(30):5910–5921, 2019.

[207] Leor N. Katz, Jacob L. Yates, Jonathan W. Pillow, and Alexander C. Huk. Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature*, 535(7611):285–288, Jul 2016. ISSN 1476-4687. doi: 10.1038/nature18617. URL https://doi.org/10.1038/nature18617.

[208] Jürgen Kayser and Craig E Tenke. Optimizing pca methodology for erp component identification and measurement: theoretical rationale and empirical evaluation. *Clinical neurophysiology*, 114(12):2307–2325, 2003.

[209] Simon P Kelly and Redmond G O'Connell. Internal and external influences on the rate of sensory evidence accumulation in the human brain. *Journal of Neuroscience*, 33(50): 19434–19441, 2013.

[210] J. A. Scott Kelso. Multistability and metastability: understanding dynamic coordination in the brain. *Philosophical transactions of the Royal Society of London. Series B,*

Biological sciences, 367(1591):906–918, Apr 2012. ISSN 1471-2970. doi: 10.1098/rstb.2011.0351. URL https://www.ncbi.nlm.nih.gov/pubmed/22371613. 22371613[pmid].

[211] Max C. Keuken, Christa Müller-Axt, Robert Langner, Simon B. Eickhoff, Birte U. Forstmann, and Jane Neumann. Brain networks of perceptual decision-making: An fMRI ale meta-analysis. *Frontiers in Human Neuroscience*, 8:445, 2014. ISSN 1662-5161. doi: 10.3389/fnhum.2014.00445. URL https://www.frontiersin.org/article/10.3389/fnhum.2014.00445.

[212] Christian Keysers, Valeria Gazzola, and Eric-Jan Wagenmakers. Using bayes factor hypothesis testing in neuroscience to establish evidence of absence. *Nature neuroscience*, 23(7):788–799, 2020.

[213] Stefan J Kiebel, Marta I Garrido, Rosalyn J Moran, and Karl J Friston. Dynamic causal modelling for EEG and MEG. *Cognitive neurodynamics*, 2(2):121, 2008.

[214] Tim C Kietzmann, Courtney J Spoerer, Lynn KA Sörensen, Radoslaw M Cichy, Olaf Hauk, and Nikolaus Kriegeskorte. Recurrence is required to capture the representational dynamics of the human visual system. *Proceedings of the National Academy of Sciences*, 116(43):21854–21863, 2019.

[215] Bommae Kim and Timo von Oertzen. Classifiers as a model-free group comparison test. *Behavior Research Methods*, 50(1):416–426, Feb 2018. ISSN 1554-3528. doi: 10.3758/s13428-017-0880-z. URL https://doi.org/10.3758/s13428-017-0880-z.

[216] Ji Hyun Kim, Jun Ki Lee, Seong-Beom Koh, Sang-Ahm Lee, Jong-Min Lee, Sun I. Kim, and Joong Koo Kang. Regional grey matter abnormalities in juvenile myoclonic epilepsy: A voxel-based morphometry study. *NeuroImage*, 37(4):1132 – 1137, 2007. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2007.06.025. URL http://www.sciencedirect.com/science/article/pii/S1053811907005423.

[217] Jong-Nam Kim and Michael N. Shadlen. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature Neuroscience*, 2(2):176–185, Feb 1999. ISSN 1546-1726. doi: 10.1038/5739. URL https://doi.org/10.1038/5739.

[218] Sun-Young Kim, Yang-Ha Hwang, Ho-Won Lee, Chung-Kyu Suh, Soon-Hak Kwon, and Sung-Pa Park. Cognitive impairment in juvenile myoclonic epilepsy. *Journal of clinical*

neurology (Seoul, Korea), 3(2):86–92, Jun 2007. ISSN 1738-6586. doi: 10.3988/jcn.2007.3.2.86. URL https://www.ncbi.nlm.nih.gov/pubmed/19513297. 19513297[pmid].

[219] Timo Kirschstein and Rüdiger Köhling. What is the source of the EEG? *Clinical EEG and neuroscience*, 40(3):146–149, 2009.

[220] Silke Klamer, Sabine Rona, Adham Elshahabi, Holger Lerche, Christoph Braun, Jürgen Honegger, Michael Erb, and Niels K Focke. Multimodal effective connectivity analysis reveals seizure focus and propagation in musicogenic epilepsy. *Neuroimage*, 113:70–77, 2015.

[221] Arno Klein and Jason Tourville. 101 labeled brain images and a consistent human cortical labeling protocol. *Frontiers in Neuroscience*, 6:171, 2012. ISSN 1662-453X. doi: 10.3389/fnins.2012.00171. URL https://www.frontiersin.org/article/10.3389/fnins.2012.00171.

[222] Mario Kleiner, David Brainard, and Denis Pelli. What's new in psychtoolbox-3? 2007.

[223] Albert Kok. On the utility of p3 amplitude as a measure of processing capacity. *Psychophysiology*, 38(3):557–577, 2001.

[224] Nils Kolling, Marco Wittmann, and Matthew FS Rushworth. Multiple neural mechanisms of decision making and their competition under changing risk pressure. *Neuron*, 81(5):1190–1202, 2014.

[225] Alexei A Koulakov, Sridhar Raghavachari, Adam Kepecs, and John E Lisman. Model for a robust neural integrator. *Nature neuroscience*, 5(8):775–782, 2002.

[226] Ian Krajbich, Dingchao Lu, Colin Camerer, and Antonio Rangel. The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in psychology*, 3:193, 2012.

[227] Ruth M Krebs, Carsten N Boehler, and Marty G Woldorff. The influence of reward associations on conflict processing in the stroop task. *Cognition*, 117(3):341–347, 2010.

[228] Nikolaus Kriegeskorte and Pamela K Douglas. Cognitive computational neuroscience. *Nature Neuroscience*, 21(9):1148–1160, Sep 2018. ISSN 1546-1726. doi: 10.1038/s41593-018-0210-5. URL https://doi.org/10.1038/s41593-018-0210-5.

[229] A Krizhevsky, I Sutskever, and GE Hinton. 2012 alexnet. *Adv. Neural Inf. Process. Syst.*, pages 1–9, 2012.

[230] Joseph Krummenacher, Hermann J Müller, and Dieter Heller. Visual search for dimensionally redundant pop-out targets: Redundancy gains in compound tasks. *Visual Cognition*, 9(7):801–837, 2002.

[231] John Kruschke. Doing bayesian data analysis: A tutorial with r, jags, and stan. 2014.

[232] Dominik Krzemiński, Sebastian Michelmann, Matthias Treder, and Lorena Santamaria. Classification of P300 component using a Riemannian ensemble approach. In *Mediterranean Conference on Medical and Biological Engineering and Computing*, pages 1885–1889. Springer, 2019.

[233] Dominik Krzemiński, Naoki Masuda, Khalid Hamandi, Krish D. Singh, Bethany Routley, and Jiaxiang Zhang. Energy landscape of resting magnetoencephalography reveals fronto-parietal network impairments in epilepsy. *Network Neuroscience*, 4(2):374–396, 2020. URL https://doi.org/10.1162/netn_a_00125.

[234] Raima Larter, Brent Speelman, and Robert M Worth. A coupled ordinary differential equation lattice model for the simulation of epileptic seizures. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 9(3):795–804, 1999.

[235] Chany Lee, Sung-Min Kim, Young-Jin Jung, Chang-Hwan Im, Dong Wook Kim, and Ki-Young Jung. Causal influence of epileptic network during spike-and-wave discharge in juvenile myoclonic epilepsy. *Epilepsy Research*, 108(2):257 – 266, 2014. ISSN 0920-1211. doi: https://doi.org/10.1016/j.eplepsyres.2013.11.005. URL http://www.sciencedirect.com/science/article/pii/S0920121113002891.

[236] Tzumin Lee, Arthur Lee, and Liqun Luo. Development of the drosophila mushroom bodies: sequential generation of three distinct types of neurons from a neuroblast. *Development*, 126(18):4065–4076, 1999.

[237] Lea Leinonen and Göte Nyman. II. functional properties of cells in anterolateral part of area 7 associative face area of awake monkeys. *Experimental Brain Research*, 34(2): 321–333, 1979.

[238] Veronika Lerche and Andreas Voss. Speed–accuracy manipulations and diffusion modeling: Lack of discriminant validity of the manipulation or of the parameter estimates? *Behavior Research Methods*, 50(6):2568–2585, Dec 2018. ISSN 1554-3528. doi: 10.3758/s13428-018-1034-7. URL https://doi.org/10.3758/s13428-018-1034-7.

[239] Daniel Levenstein, Veronica Alvarez, Asohan Amarasingham, Habiba Azab, Richard Gerkin, Andrea Hasenstaub, Ramakrishnan Iyer, Renaud Jolivet, Sarah Marzen, Joseph Monaco, Astrid Prinz, Salma Quraishi, Fidel Santamaria, Sabyasachi Shivkumar, Matthew Singh, David Stockton, Roger Traub, Horacio Rotstein, Farzan Nadim, and A. Redish. On the role of theory and modeling in neuroscience. *arXiv - preprint*, 03 2020.

[240] H. Levitt. Transformed up-down methods in psychoacoustics. *The Journal of the Acoustical Society of America*, 49(2B):467–477, 1971. doi: 10.1121/1.1912375. URL https://doi.org/10.1121/1.1912375.

[241] Qifu Li, Weifang Cao, Xiaoping Liao, Zhibin Chen, Tianhua Yang, Qiyong Gong, Dong Zhou, Cheng Luo, and Dezhong Yao. Altered resting state functional network connectivity in children absence epilepsy. *Journal of the Neurological Sciences*, 354(1): 79 – 85, 2015. ISSN 0022-510X. doi: https://doi.org/10.1016/j.jns.2015.04.054. URL http://www.sciencedirect.com/science/article/pii/S0022510X15002658.

[242] Yuxi Li. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*, 2017.

[243] Wei Liao, Zhiqiang Zhang, Zhengyong Pan, Dante Mantini, Jurong Ding, Xujun Duan, Cheng Luo, Guangming Lu, and Huafu Chen. Altered functional connectivity and small-world in mesial temporal lobe epilepsy. *PLOS ONE*, 5(1):1–11, 01 2010. doi: 10.1371/journal.pone.0008525. URL https://doi.org/10.1371/journal.pone.0008525.

[244] Wei Liao, Zhiqiang Zhang, Zhengyong Pan, Dante Mantini, Jurong Ding, Xujun Duan, Cheng Luo, Zhengge Wang, Qifu Tan, Guangming Lu, and Huafu Chen. Default mode network abnormalities in mesial temporal lobe epilepsy: A study combining fMRI and DTI. *Human Brain Mapping*, 32(6):883–895, 2011. doi: 10.1002/hbm.21076. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.21076.

[245] Chiara Della Libera and Leonardo Chelazzi. Visual selective attention and the effects of monetary rewards. *Psychological science*, 17(3):222–227, 2006.

[246] Charles C Liu and Takeo Watanabe. Accounting for speed–accuracy tradeoff in perceptual learning. *Vision research*, 61:107–114, 2012.

[247] Marinho A Lopes, Marc Goodfellow, and John R Terry. A model-based assessment of the seizure onset zone predictive power to inform the epileptogenic zone. *Frontiers in computational neuroscience*, 13:25, 2019.

[248] Marinho A Lopes, Dominik Krzemiński, Khalid Hamandi, Krish D Singh, Naoki Masuda, John R Terry, and Jiaxiang Zhang. A computational biomarker of juvenile myoclonic epilepsy from resting-state MEG. *Clinical Neurophysiology*, 2020. doi: https://doi.org/10.1016/j.clinph.2020.12.021.

[249] Gerard M Loughnane, Daniel P Newman, Mark A Bellgrove, Edmund C Lalor, Simon P Kelly, and Redmond G O'Connell. Target selection signals influence perceptual decisions by modulating the onset and rate of evidence accumulation. *Current Biology*, 26(4): 496–502, 2016.

[250] Bradley C Love. Levels of biological plausibility. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 376(1815):20190632, January 2021. ISSN 0962-8436. doi: 10.1098/rstb.2019.0632. URL https://doi.org/10.1098/rstb.2019.0632.

[251] Steven J Luck and Steven A Hillyard. Spatial filtering during visual search: evidence from human electrophysiology. *Journal of Experimental Psychology: Human Perception and Performance*, 20(5):1000, 1994.

[252] Kitty K Lui, Michael D Nunez, Jessica M Cassidy, Joachim Vandekerckhove, Steven C Cramer, and Ramesh Srinivasan. Timing of readiness potentials reflect a decision-making process in the human brain. *bioRxiv*, page 338806, 2018.

[253] John G Lynch, JW Alba, and J Wesley Hutchinson. Memory and decision making. *Handbook of Consumer Behavior*, pages 1–9, 1991.

[254] Jakob H Macke, Iain Murray, and Peter E Latham. How biased are maximum entropy models? *Advances in Neural Information Processing Systems*, pages 1–9, 2012.

[255] Leah Maizey and Loukia Tzavella. Barriers and solutions for early career researchers in tackling the reproducibility crisis in cognitive neuroscience. *Cortex*, 113:357–359, 2019.

[256] George R Mangun and Steven A Hillyard. Modulations of sensory-evoked brain potentials indicate changes in perceptual processing during visual-spatial priming. *Journal of Experimental Psychology: Human perception and performance*, 17(4):1057, 1991.

[257] Uri Maoz, Gideon Yaffe, Christof Koch, and Liad Mudrik. Neural precursors of decisions that matter—an erp study of deliberate and arbitrary choice. *eLife*, 8, 2019.

[258] Eric Maris and Robert Oostenveld. Nonparametric statistical testing of EEG-and MEG-data. *Journal of neuroscience methods*, 164(1):177–190, 2007.

[259] Henry Markram. The Blue Brain Project. *Nature Reviews Neuroscience*, 7(2):153–160, Feb 2006. ISSN 1471-0048. doi: 10.1038/nrn1848. URL https://doi.org/10.1038/nrn1848.

[260] Henry Markram, Karlheinz Meier, Thomas Lippert, Sten Grillner, Richard Frackowiak, Stanislas Dehaene, Alois Knoll, Haim Sompolinsky, Kris Verstreken, Javier DeFelipe, Seth Grant, Jean-Pierre Changeux, and Alois Saria. Introducing the human brain project. *Procedia Computer Science*, 7:39 – 42, 2011. ISSN 1877-0509. doi: https://doi.org/10.1016/j.procs.2011.12.015. URL http://www.sciencedirect.com/science/article/pii/S1877050911006806. Proceedings of the 2nd European Future Technologies Conference and Exhibition 2011 (FET 11).

[261] David Marr. *Vision: A computational investigation into the human representation and processing of visual information*. W.H. Freeman, New York, NY, 1983.

[262] James AR Marshall, Rafal Bogacz, and Iain D Gilchrist. Consistent implementation of decisions in the brain. *PloS one*, 7(9), 2012.

[263] Hannes Matuschek, Reinhold Kliegl, Shravan Vasishth, Harald Baayen, and Douglas Bates. Balancing type i error and power in linear mixed models. *Journal of Memory and Language*, 94:305–315, 2017.

[264] Mark E. Mazurek, Jamie D. Roitman, Jochen Ditterich, and Michael N. Shadlen. A Role for Neural Integrators in Perceptual Decision Making. *Cerebral Cortex*, 13(11):

1257–1269, 11 2003. ISSN 1047-3211. doi: 10.1093/cercor/bhg097. URL https://doi.org/10.1093/cercor/bhg097.

[265] James L McClelland, David E Rumelhart, PDP Research Group, et al. Parallel distributed processing. *Explorations in the Microstructure of Cognition*, 2:216–271, 1986.

[266] Megan L. McGill, Orrin Devinsky, Clare Kelly, Michael Milham, F. Xavier Castellanos, Brian T. Quinn, Jonathan DuBois, Jonathan R. Young, Chad Carlson, Jacqueline French, Ruben Kuzniecky, Eric Halgren, and Thomas Thesen. Default mode network abnormalities in idiopathic generalized epilepsy. *Epilepsy & Behavior*, 23(3):353 – 359, 2012. ISSN 1525-5050. doi: https://doi.org/10.1016/j.yebeh.2012.01.013. URL http://www.sciencedirect.com/science/article/pii/S1525505012000285.

[267] Christoph M. Michel and Thomas Koenig. EEG microstates as a tool for studying the temporal dynamics of whole-brain neuronal networks: A review. *NeuroImage*, 180:577–593, 2018. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2017.11.062. URL http://www.sciencedirect.com/science/article/pii/S105381191731008X.

[268] Ewa A Miendlarzewska, Daphne Bavelier, and Sophie Schwartz. Influence of reward motivation on human declarative memory. *Neuroscience & Biobehavioral Reviews*, 61:156–176, 2016.

[269] Steven Miletić and Leendert van Maanen. Caution in decision-making under time pressure is mediated by timing ability. *Cognitive Psychology*, 110:16–29, 2019. ISSN 0010-0285(Print). doi: 10.1016/j.cogpsych.2019.01.002. URL https://doi.org/10.1016/j.cogpsych.2019.01.002.

[270] Ludovico Minati, Giulia Varotto, Ludovico D'Incerti, Ferruccio Panzica, and Dennis Chan. From brain topography to brain topology: relevance of graph theory to functional neuroscience. *Neuroreport*, 24(10):536–543, 2013.

[271] Marvin Minsky. *Society of Mind*. Simon and Schuster, 1988.

[272] Rosalyn J Moran, Dimitris A Pinotsis, and Karl J Friston. Neural masses and fields in dynamic causal modeling. *Frontiers in computational neuroscience*, 7:57, 2013.

[273] VB Mountcastle, MA Steinmetz, and R Romo. Frequency discrimination in the sense of flutter: Psychophysical measurements correlated with postcentral events in behaving monkeys. *Journal of Neuroscience*, 10(9):3032–3044, 1990. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.10-09-03032.1990. URL https://www.jneurosci.org/content/10/9/3032.

[274] Martijn J Mulder, Eric-Jan Wagenmakers, Roger Ratcliff, Wouter Boekel, and Birte U Forstmann. Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *Journal of Neuroscience*, 32(7):2335–2343, 2012.

[275] Martijn J Mulder, Max C Keuken, Leendert van Maanen, Wouter Boekel, Birte U Forstmann, and Eric-Jan Wagenmakers. The speed and accuracy of perceptual decisions in a random-tone pitch task. *Attention, Perception, & Psychophysics*, 75(5):1048–1058, 2013.

[276] MJ Mulder, L Van Maanen, and BU Forstmann. Perceptual decision neurosciences–a model-based review. *Neuroscience*, 277:872–884, 2014.

[277] Philip Munz, Ioan Hudea, Joe Imad, and Robert J Smith? When zombies attack!: Mathematical modelling of an outbreak of zombie infection. *Infectious Disease modelling Research Progress*, 4:133–150, 2009.

[278] Suresh D. Muthukumaraswamy, Robin L. Carhart-Harris, Rosalyn J. Moran, Matthew J. Brookes, Tim M. Williams, David Errtizoe, Ben Sessa, Andreas Papadopoulos, Mark Bolstridge, Krish D. Singh, Amanda Feilding, Karl J. Friston, and David J. Nutt. Broadband cortical desynchronization underlies the human psychedelic state. *Journal of Neuroscience*, 33(38):15171–15183, 2013. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.2063-13.2013. URL http://www.jneurosci.org/content/33/38/15171.

[279] Allen Newell. You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. *Technical Report*, 1973.

[280] Allen Newell, Herbert Alexander Simon, et al. *Human Problem Solving*, volume 104. Prentice-Hall Englewood Cliffs, NJ, 1972.

[281] William T Newsome, Kenneth H Britten, and J Anthony Movshon. Neuronal correlates of a perceptual decision. *Nature*, 341(6237):52–54, 1989.

[282] WT Newsome and EB Pare. A selective impairment of motion perception following lesions of the middle temporal visual area (mt). *Journal of Neuroscience*, 8(6):2201–2211, 1988. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.08-06-02201.1988. URL https://www.jneurosci.org/content/8/6/2201.

[283] Lopes da Silva-F. H. Niedermeyer, Ernst. *Electroencephalography*. Lippincott Williams & Wilkins, 2005.

[284] Sander Nieuwenhuis, Gary Aston-Jones, and Jonathan D Cohen. Decision making, the P3, and the locus coeruleus–norepinephrine system. *Psychological bulletin*, 131(4):510, 2005.

[285] Guiomar Niso, Sira Carrasco, María Gudín, Fernando Maestú, Francisco del Pozo, and Ernesto Pereda. What graph theory actually tells us about resting state interictal MEG epileptic activity. *NeuroImage: Clinical*, 8:503 – 515, 2015. ISSN 2213-1582. doi: https://doi.org/10.1016/j.nicl.2015.05.008. URL http://www.sciencedirect.com/science/article/pii/S221315821500100X.

[286] Ritwik K Niyogi and KongFatt Wong-Lin. Dynamic excitatory and inhibitory gain modulation can produce flexible, robust and optimal decision-making. *PLoS Comput Biol*, 9(6):e1003099, 2013.

[287] Guido Nolte. The magnetic lead field theorem in the quasi-static approximation and its use for magnetoencephalography forward calculation in realistic volume conductors. *Physics in Medicine and Biology*, 48(22):3637–3652, oct 2003. doi: 10.1088/0031-9155/48/22/002. URL https://doi.org/10.1088%2F0031-9155%2F48%2F22%2F002.

[288] Michael D Nunez, Joachim Vandekerckhove, and Ramesh Srinivasan. How attention influences perceptual decision making: Single-trial EEG correlates of drift-diffusion model parameters. *Journal of mathematical psychology*, 76:117–130, 2017.

[289] Michael D Nunez, Aishwarya Gosai, Joachim Vandekerckhove, and Ramesh Srinivasan. The latency of a visual evoked potential tracks the onset of decision making. *Neuroimage*, 197:93–108, 2019.

[290] P.L. Nunez and R.D. Katznelson. *Electric Fields of the Brain: The Neurophysics of EEG*. Oxford Univ. Press, 1981. ISBN 9780195027969. URL https://books.google.co.uk/books?id=gu5qAAAAMAAJ.

[291] Hanna Oh-Descher, Jeffrey M Beck, Silvia Ferrari, Marc A Sommer, and Tobias Egner. Probabilistic inference under time pressure leads to a cortical-to-subcortical shift in decision evidence integration. *NeuroImage*, 162, 2017.

[292] Andrés Ojeda, Robin A Murphy, and Alex Kacelnik. Paradoxical choice in rats: Subjective valuation and mechanism of choice. *Behavioural processes*, 152:73–80, 2018.

[293] Piotr Olejniczak. Neurophysiologic basis of EEG. *Journal of clinical neurophysiology*, 23(3):186–189, 2006.

[294] J. O'Muircheartaigh, C. Vollmar, G.J. Barker, V. Kumari, M.R. Symms, P. Thompson, J.S. Duncan, M.J. Koepp, and M.P. Richardson. Focal structural changes and cognitive dysfunction in juvenile myoclonic epilepsy. *Neurology*, 76(1):34–40, 2011. ISSN 0028-3878. doi: 10.1212/WNL.0b013e318203e93d. URL https://n.neurology.org/content/76/1/34.

[295] Adam F Osth, Simon Dennis, and Andrew Heathcote. Likelihood ratio sequential sampling models of recognition memory. *Cognitive Psychology*, 92:101–126, 2017.

[296] Bastiaan Oud, Ian Krajbich, Kevin Miller, Jin Hyun Cheong, Matthew Botvinick, and Ernst Fehr. Irrational time allocation in decision-making. *Proceedings of the Royal Society B: Biological Sciences*, 283(1822):20151439, 2016.

[297] Redmond G O'Connell, Michael N Shadlen, KongFatt Wong-Lin, and Simon P Kelly. Bridging neural and computational viewpoints on perceptual decision-making. *Trends in neurosciences*, 41(11):838–852, 2018.

[298] Darren Pais, Patrick M Hogan, Thomas Schlegel, Nigel R Franks, Naomi E Leonard, and James AR Marshall. A mechanism for value-sensitive decision-making. *PloS one*, 8(9), 2013.

[299] John Palmer. Attention in visual search: Distinguishing four causes of a set-size effect. *Current directions in psychological science*, 4(4):118–123, 1995.

[300] John Palmer, Alexander C Huk, and Michael N Shadlen. The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of vision*, 5(5):1–1, 2005.

[301] Thomas J Palmeri, Jeffrey D Schall, and Gordon D Logan. Neurocognitive modeling of perceptual decision making. 2015.

[302] A.C. Papanicolaou, P. Pazo-Alvarez, E.M. Castillo, R.L. Billingsley-Marshall, J.I. Breier, P.R. Swank, S. Buchanan, M. McManis, T. Clear, and A.D. Passaro. Functional neuroimaging with MEG: Normative language profiles. *NeuroImage*, 33(1):326 – 342, 2006. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2006.06.020. URL http://www.sciencedirect.com/science/article/pii/S1053811906006896.

[303] Lucas C Parra, Clay D Spence, Adam D Gerson, and Paul Sajda. Recipes for the linear analysis of EEG. *Neuroimage*, 28(2):326–341, 2005.

[304] Roberto D Pascual-Marqui, Michaela Esslen, Kieko Kochi, Dietrich Lehmann, et al. Functional imaging with low-resolution brain electromagnetic tomography (loreta): a review. *Methods and findings in experimental and clinical pharmacology*, 24(Suppl C): 91–95, 2002.

[305] HN Phillips, TE Cope, LE Hughes, Jiaxiang Zhang, and JB Rowe. Monitoring the past and choosing the future: the prefrontal cortical influences on voluntary action. *Scientific reports*, 8(1):1–13, 2018.

[306] Ada Piazzini, Katherine Turner, Aglaia Vignoli, Raffaele Canger, and Maria Paola Canevini. Frontal cognitive dysfunction in juvenile myoclonic epilepsy. *Epilepsia*, 49(4): 657–662, 2008. doi: 10.1111/j.1528-1167.2007.01482.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1528-1167.2007.01482.x.

[307] Angelo Pirrone, Habiba Azab, Benjamin Y Hayden, Tom Stafford, and James AR Marshall. Evidence for the speed–value trade-off: Human and monkey decision making is magnitude sensitive. *Decision*, 5(2):129, 2018.

[308] M Andrea Pisauro, Elsa Fouragnan, Chris Retzler, and Marios G Philiastides. Neural correlates of evidence accumulation during value-based decisions revealed via simultaneous EEG-fMRI. *Nature Communications*, 8(1):1–9, 2017.

[309] Rüdiger F Pohl. On the use of recognition in inferential decision making: An overview of the debate. *Judgment and Decision Making*, 6(5):423, 2011.

[310] Rafael Polanía, Ian Krajbich, Marcus Grueschow, and Christian C Ruff. Neural oscillations and synchronization differentially support evidence accumulation in perceptual and value-based decision making. *Neuron*, 82(3):709–720, 2014.

[311] John Polich, Patricia Crane Ellerson, and Jill Cohen. P300, stimulus intensity, modality, and probability. *International Journal of Psychophysiology*, 23(1-2):55–62, 1996.

[312] Jean-Baptiste Poline and Matthew Brett. The general linear model and fMRI: does love last forever? *Neuroimage*, 62(2):871–880, 2012.

[313] Cambridge Univ. Press. *Cambridge Advanced Learner's Dictionary*. Cambridge Univ. Press, 2008. ISBN 9780521674683.

[314] Andrew J. Quinn, Diego Vidaurre, Romesh Abeysuriya, Robert Becker, Anna C. Nobre, and Mark W. Woolrich. Task-evoked dynamic network analysis through hidden markov modeling. *Frontiers in Neuroscience*, 12:603, 2018. ISSN 1662-453X. doi: 10.3389/fnins.2018.00603. URL https://www.frontiersin.org/article/10.3389/fnins.2018.00603.

[315] Rodrigo Quian Quiroga. Concept cells: the building blocks of declarative memory functions. *Nature Reviews Neuroscience*, 13(8):587–597, 2012.

[316] Sivananda Rajananda, Hakwan Lau, and Brian Odegaard. A random-dot kinematogram for web-based vision research. *J. Open Res. Softw.*, 6, 2018.

[317] Antonio Rangel and John A Clithero. The computation of stimulus values in simple choice. In *Neuroeconomics*. Elsevier, 2014.

[318] Antonio Rangel, Colin Camerer, and P. Read Montague. A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7): 545–556, 2008. doi: 10.1038/nrn2357. URL https://doi.org/10.1038/nrn2357.

[319] Dragan Rangelov and Jason B Mattingley. Evidence accumulation during perceptual decision-making is sensitive to the dynamics of attentional selection. *NeuroImage*, 220: 117093, 2020.

[320] Rajesh PN Rao and Dana H Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1): 79–87, 1999.

[321] Roger Ratcliff. A theory of memory retrieval. *Psychological review*, 85(2):59, 1978.

[322] Roger Ratcliff. A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin & Review*, 9(2):278–291, Jun 2002. ISSN 1531-5320. doi: 10.3758/BF03196283. URL https://doi.org/10.3758/BF03196283.

[323] Roger Ratcliff. Modeling response signal and response time data. *Cognitive psychology*, 53(3):195–237, 2006.

[324] Roger Ratcliff and Michael J Frank. Reinforcement-based decision making in corticostriatal circuits: mutual constraints by neurocomputational and diffusion models. *Neural computation*, 24(5):1186–1229, 2012.

[325] Roger Ratcliff and Gail McKoon. The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*, 20(4):873–922, 2008.

[326] Roger Ratcliff and Jeffrey N Rouder. A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human perception and performance*, 26(1):127, 2000.

[327] Roger Ratcliff and Philip L Smith. A comparison of sequential sampling models for two-choice reaction time. *Psychological review*, 111(2):333, 2004.

[328] Roger Ratcliff and Francis Tuerlinckx. Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic bulletin & review*, 9(3):438–481, 2002.

[329] Roger Ratcliff, Chelsea Voskuilen, and Andrei Teodorescu. Modeling 2-alternative forced-choice tasks: Accounting for both magnitude and difference effects. *Cognitive psychology*, 103:1–22, 2018.

[330] Jane E Raymond and Jennifer L O'Brien. Selective visual attention and motivation: The consequences of value learning in an attentional blink task. *Psychological Science*, 20(8):981–988, 2009.

[331] John H. Reynolds and Leonardo Chelazzi. Attentional modulation of visual processing. *Annual Review of Neuroscience*, 27(1):611–647, 2004. doi: 10.1146/annurev.neuro.26.041002.131039. URL https://doi.org/10.1146/annurev.neuro.26.041002.131039. PMID: 15217345.

[332] Danilo Rezende, Ivo Danihelka, Karol Gregor, Daan Wierstra, et al. One-shot generalization in deep generative models. In *International Conference on Machine Learning*, pages 1521–1529. PMLR, 2016.

[333] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 'why should i trust you?' explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144, 2016.

[334] Fred Rieke. *Spikes: exploring the neural code*. MIT press, 1999.

[335] Jörg Rieskamp and Ulrich Hoffrage. Inferences under time pressure: How opportunity costs affect strategy selection. *Acta Psychologica*, 127(2), 2008.

[336] Jamie D Roitman and Michael N Shadlen. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of neuroscience*, 22(21):9475–9489, 2002.

[337] Edmund T. Rolls, Marc Joliot, and Nathalie Tzourio-Mazoyer. Implementation of a new parcellation of the orbitofrontal cortex in the automated anatomical labeling atlas. *NeuroImage*, 122:1 – 5, 2015. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2015.07.075. URL http://www.sciencedirect.com/science/article/pii/S1053811915006953.

[338] Lisa Ronan, Saud Alhusaini, Cathy Scanlon, Colin P. Doherty, Norman Delanty, and Mary Fitzsimons. Widespread cortical morphologic changes in juvenile myoclonic epilepsy: Evidence from structural mri. *Epilepsia*, 53(4):651–658, 2012. doi: 10.

1111/j.1528-1167.2012.03413.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1528-1167.2012.03413.x.

[339] Cristina Rosazza and Ludovico Minati. Resting-state brain networks: literature review and clinical applications. *Neurological Sciences*, 32(5):773–785, Oct 2011. ISSN 1590-3478. doi: 10.1007/s10072-011-0636-y. URL https://doi.org/10.1007/s10072-011-0636-y.

[340] Yasser Roudi, Sheila Nirenberg, and Peter E Latham. Pairwise maximum entropy models for studying large biological systems: when they can work and when they can't. *PLoS Comput Biol*, 5(5):e1000380, 2009.

[341] Bethany C Routley, Krish D Singh, Khalid Hamandi, and Suresh D Muthukumaraswamy. The effects of ampa receptor blockade on resting magnetoencephalography recordings. *Journal of Psychopharmacology*, 31(12):1527–1536, 2017. doi: 10.1177/0269881117736915. URL https://doi.org/10.1177/0269881117736915. PMID: 29084475.

[342] Aldo Rustichini. Neuroeconomics: Formal models of decision making and cognitive neuroscience. In *Neuroeconomics*, pages 33–46. Elsevier, 2009.

[343] Raymond Salvador, John Suckling, Martin R. Coleman, John D. Pickard, David Menon, and Ed Bullmore. Neurophysiological Architecture of Functional Magnetic Resonance Images of Human Brain. *Cerebral Cortex*, 15(9):1332–1342, 01 2005. ISSN 1047-3211. doi: 10.1093/cercor/bhi016. URL https://doi.org/10.1093/cercor/bhi016.

[344] M Salvaris and F Sepulveda. Visual modifications on the p300 speller BCI paradigm. *Journal of Neural Engineering*, 6(4):046011, jul 2009. doi: 10.1088/1741-2560/6/4/046011. URL https://doi.org/10.1088/1741-2560/6/4/046011.

[345] Arthur L Samuel. Some studies in machine learning using the game of checkers. II—recent progress. *IBM Journal of Research and Development*, 11(6):601–617, 1967.

[346] Jose M Sanchez-Bornot, Maria E Lopez, Ricardo Bruña, Fernando Maestu, Vahab Youssofzadeh, Su Yang, David P Finn, Stephen Todd, Paula L McLean, Girijesh Prasad, et al. High-dimensional brain-wide functional connectivity mapping in magnetoencephalography. *Journal of Neuroscience Methods*, 348:108991, 2021.

[347] Paula Sanz Leon, Stuart Knock, M. Woodman, Lia Domide, Jochen Mersmann, Anthony McIntosh, and Viktor Jirsa. The virtual brain: a simulator of primate brain network dynamics. *Frontiers in Neuroinformatics*, 7:10, 2013. ISSN 1662-5196. doi: 10.3389/fninf.2013.00010. URL https://www.frontiersin.org/article/10.3389/fninf.2013.00010.

[348] J Sarnthein, A Morel, A Von Stein, and D Jeanmonod. Thalamic theta field potentials and EEG: high thalamocortical coherence in patients with neurogenic pain, epilepsy and movement disorders. *Thalamus & Related Systems*, 2(3):231–238, 2003.

[349] T. Sarwar, Y. Tian, B.T.T. Yeo, K. Ramamohanarao, and A. Zalesky. Structure-function coupling in the human connectome: A machine learning approach. *NeuroImage*, 226:117609, 2021. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2020.117609. URL http://www.sciencedirect.com/science/article/pii/S1053811920310946.

[350] I. Savic, A. Lekvall, D. Greitz, and G. Helms. Mr spectroscopy shows reduced frontal lobe concentrations of n-acetyl aspartate in patients with juvenile myoclonic epilepsy. *Epilepsia*, 41(3):290–296, 2000. doi: 10.1111/j.1528-1157.2000.tb00158.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1528-1157.2000.tb00158.x.

[351] Jeffrey D Schall. Decision making: From sensory evidence to a motor command. *Current Biology*, 10(11):R404–R406, 2000.

[352] Johannes Schiebener and Matthias Brand. Decision making under objective risk conditions–a review of cognitive and emotional correlates, strategies, feedback processing, and external influences. *Neuropsychology review*, 25(2):171–198, 2015.

[353] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.

[354] Florian Schmiedek, Klaus Oberauer, Oliver Wilhelm, Heinz-Martin Süß, and Werner W Wittmann. Individual differences in components of reaction time distributions and their relations to working memory and intelligence. *Journal of Experimental Psychology: General*, 136(3):414, 2007.

[355] Matthew T Schmolesky, Youngchang Wang, Doug P Hanes, Kirk G Thompson, Stefan Leutgeb, Jeffrey D Schall, and Audie G Leventhal. Signal timing across the macaque visual system. *Journal of neurophysiology*, 79(6):3272–3278, 1998.

[356] Elad Schneidman, Michael J. Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006. ISSN 1476-4687. doi: 10.1038/nature04701. URL https://doi.org/10.1038/nature04701.

[357] Jacqueline Scholl, Nils Kolling, Natalie Nelissen, Marco K Wittmann, Catherine J Harmer, and Matthew FS Rushworth. The good, the bad, and the irrelevant: neural mechanisms of learning real and hypothetical rewards and effort. *Journal of Neuroscience*, 35(32):11233–11251, 2015.

[358] Nicolas W Schuck, Robert Gaschler, Dorit Wenke, Jakob Heinzle, Peter A Frensch, John-Dylan Haynes, and Carlo Reverberi. Medial prefrontal cortex predicts internally driven strategy shifts. *Neuron*, 86(1):331–340, 2015.

[359] Wolfram Schultz. Behavioral dopamine signals. *Trends in neurosciences*, 30(5):203–210, 2007.

[360] DL Schurman and JP Belcher. Noncontingent reward magnitude effects on reaction time: A replication and extension. *Bulletin of the Psychonomic Society*, 3(2):104–106, 1974.

[361] Kilian Semmelmann and Sarah Weigelt. Online psychophysics: Reaction time effects in cognitive experiments. *Behavior Research Methods*, 49(4):1241–1260, 2017.

[362] Michael N Shadlen and Roozbeh Kiani. Decision making as a window on cognition. *Neuron*, 80(3):791–806, 2013.

[363] Michael N Shadlen and Daphna Shohamy. Decision making and sequential sampling from memory. *Neuron*, 90(5):927–939, 2016.

[364] MN Shadlen, KH Britten, WT Newsome, and JA Movshon. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *Journal of Neuroscience*, 16(4):1486–1510, 1996. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.16-04-01486.1996. URL https://www.jneurosci.org/content/16/4/1486.

[365] Marilyn L Shaw. Attending to multiple sources of information: I. the integration of information in decision making. *Cognitive Psychology*, 14(3):353 – 409, 1982. ISSN 0010-0285. doi: https://doi.org/10.1016/0010-0285(82)90014-7. URL http://www.sciencedirect.com/science/article/pii/0010028582900147.

[366] Galit Shmueli. To explain or to predict? *Statist. Sci.*, 25(3):289–310, 08 2010. doi: 10.1214/10-STS330. URL https://doi.org/10.1214/10-STS330.

[367] Shepard Siegel and Lorraine G Allan. The widespread influence of the rescorla-wagner model. *Psychonomic Bulletin & Review*, 3(3):314–321, 1996.

[368] Ray Singh and Chris Eliasmith. Higher-dimensional neurons explain the tuning and dynamics of working memory cells. *The Journal of Neuroscience*, 26(14), 2006.

[369] Yuedong Song and Jiaxiang Zhang. Automatic recognition of epileptic EEG patterns via extreme learning machine and multiresolution feature extraction. *Expert Systems with Applications*, 40(14):5477 – 5489, 2013. ISSN 0957-4174. doi: https://doi.org/10.1016/j.eswa.2013.04.025. URL http://www.sciencedirect.com/science/article/pii/S0957417413002625.

[370] Alexander Sorokin and David Forsyth. Utility data annotation with amazon mechanical turk. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 1–8. IEEE, 2008.

[371] David J Spiegelhalter, Nicola G Best, Bradley P Carlin, and Angelika Van Der Linde. Bayesian measures of model complexity and fit. *Journal of the royal statistical society: Series b (statistical methodology)*, 64(4):583–639, 2002.

[372] C. J. Stam and E. C. W. van Straaten. The organization of physiological brain networks. *Clinical Neurophysiology*, 123(6):1067 – 1087, 2012. ISSN 1388-2457. doi: https://doi.org/10.1016/j.clinph.2012.01.011. URL http://www.sciencedirect.com/science/article/pii/S1388245712000570.

[373] C.J. Stam. Nonlinear dynamical analysis of EEG and MEG: Review of an emerging field. *Clinical Neurophysiology*, 116(10):2266 – 2301, 2005. ISSN 1388-2457. doi: https://doi.org/10.1016/j.clinph.2005.06.011. URL http://www.sciencedirect.com/science/article/pii/S1388245705002403.

[374] Dominic Standage, Gunnar Blohm, and Michael C. Dorris. On the neural implementation of the speed-accuracy trade-off. *Frontiers in Neuroscience*, 8:236, 2014. ISSN 1662-453X. doi: 10.3389/fnins.2014.00236. URL https://www.frontiersin.org/article/10.3389/fnins.2014.00236.

[375] Dominic Standage, Da-Hui Wang, and Gunnar Blohm. Neural dynamics implement a flexible decision bound with a fixed firing rate for choice: a model-based hypothesis. *Frontiers in Neuroscience*, 8:318, 2014. ISSN 1662-453X. doi: 10.3389/fnins.2014.00318. URL https://www.frontiersin.org/article/10.3389/fnins.2014.00318.

[376] Jeffrey J Starns and Roger Ratcliff. Validating the unequal-variance assumption in recognition memory using response time distributions instead of roc functions: A diffusion model analysis. *Journal of memory and language*, 70:36–52, 2014.

[377] Leon Stefanovski, Jil Mona Meier, Roopa Kalsank Pai, Paul Triebkorn, Tristram Lett, Leon Martin, Konstantin Bülau, Martin Hofmann-Apitius, Ana Solodkin, Anthony Randal McIntosh, et al. From genes to network models of alzheimer's disease: Biological framework for multi-scale brain simulation with the virtual brain. 2020.

[378] Terrence Stewart and Chris Eliasmith. Compositionality and biologically plausible models. 2012.

[379] Terrence C Stewart and Chris Eliasmith. Neural cognitive modelling: A biologically constrained spiking neuron model of the tower of hanoi task. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, 2011.

[380] Terrence C Stewart, Xuan Choo, and Chris Eliasmith. Dynamic behaviour of a spiking model of action selection in the basal ganglia. In *Proceedings of the 10th International Conference on Cognitive Modeling*. Philadelphia, USA, 2010.

[381] & Eliasmith C. Stewart T., Choo F. Spaun: A perception-cognition-action model using spiking neurons. 34. Annual Meeting of the Cognitive Science Society, 2012. URL https://escholarship.org/uc/item/168466tf.

[382] Marcel Stimberg, Romain Brette, and Dan FM Goodman. Brian 2, an intuitive and efficient neural simulator. *Elife*, 8:e47314, 2019.

[383] Katherine R. Storrs and Guido Maiello. A model for neural network modeling in neuroscience. *Journal of Neuroscience*, 40(37):7010–7012, 2020. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.1205-20.2020. URL https://www.jneurosci.org/content/40/37/7010.

[384] Steven H Strogatz. *Nonlinear dynamics and chaos with student solutions manual: With applications to physics, biology, chemistry, and engineering*. CRC press, 2018.

[385] Irene Sturm, Sebastian Lapuschkin, Wojciech Samek, and Klaus-Robert Müller. Interpretable deep neural networks for single-trial EEG classification. *Journal of neuroscience methods*, 274:141–145, 2016.

[386] Christopher Summerfield and Konstantinos Tsetsos. Building bridges between perceptual and economic decision-making: neural and computational mechanisms. *Frontiers in neuroscience*, 6:70, 2012.

[387] Ron Sun and Frederic Alexandre. *Connectionist-symbolic integration: From unified to hybrid approaches*. Psychology Press, 2013.

[388] B.E Swartz, E Halgren, F Simpkins, and K Syndulko. Primary memory in patients with frontal and primary generalized epilepsy. *Journal of Epilepsy*, 7(3):232 – 241, 1994. ISSN 0896-6974. doi: https://doi.org/10.1016/0896-6974(94)90034-5. URL http://www.sciencedirect.com/science/article/pii/0896697494900345.

[389] Takamitsu Watanabe Takahiro Ezaki, Masayuki Ohzeki, and Naoki Masuda. Energy landscape analysis of neuroimaging data. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2096):20160287, 2017. doi: 10.1098/rsta.2016.0287. URL https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2016.0287.

[390] Aonan Tang, David Jackson, Jon Hobbs, Wei Chen, Jodi L. Smith, Hema Patel, Anita Prieto, Dumitru Petrusca, Matthew I. Grivich, Alexander Sher, Pawel Hottowy, Wladyslaw Dabrowski, Alan M. Litke, and John M. Beggs. A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. *Journal of Neuroscience*, 28(2):505–518, 2008. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.3359-07.2008. URL http://www.jneurosci.org/content/28/2/505.

[391] Andrei R Teodorescu, Rani Moran, and Marius Usher. Absolutely relative or relatively absolute: violations of value invariance in human decision making. *Psychonomic bulletin & review*, 23(1):22–38, 2016.

[392] Prejaas Tewarie, Menno M. Schoonheim, Cornelis J. Stam, Marieke L. van der Meer, Bob W. van Dijk, Frederik Barkhof, Chris H. Polman, and Arjan Hillebrand. Cognitive and clinical dysfunction, altered MEG resting-state networks and thalamic atrophy in multiple sclerosis. *PLOS ONE*, 8(7):1–11, 07 2013. doi: 10.1371/journal.pone.0069318. URL https://doi.org/10.1371/journal.pone.0069318.

[393] Richard H Thaler. Some empirical evidence on dynamic inconsistency. *Quasi rational economics*, 1:127–136, 1991.

[394] Julie Thomas, Giovanna Vanni-Mercier, and Jean-Claude Dreher. Neural dynamics of reward probability coding: a magnetoencephalographic study in humans. *Frontiers in neuroscience*, 7:214, 2013.

[395] David Thura and Paul Cisek. The basal ganglia do not select reach targets but control the urgency of commitment. *Neuron*, 95(5):1160–1170, 2017.

[396] David Thura, Julie Beauregard-Racine, Charles-William Fradet, and Paul Cisek. Decision making by urgency gating: theory and experimental support. *Journal of neurophysiology*, 108(11):2912–2930, 2012.

[397] Tim M Tierney, Niall Holmes, Stephanie Mellor, José David López, Gillian Roberts, Ryan M Hill, Elena Boto, James Leggett, Vishal Shah, Matthew J Brookes, et al. Optically pumped magnetometers: From quantum origins to multi-channel magnetoencephalography. *NeuroImage*, 199:598–608, 2019.

[398] Gabriel Tillman, Adam F Osth, Don van Ravenzwaaij, and Andrew Heathcote. A diffusion decision model analysis of evidence variability in the lexical decision task. *Psychonomic bulletin & review*, 24(6):1949–1956, 2017.

[399] Gabriel Tillman, Trish Van Zandt, and Gordon D Logan. Sequential sampling models without random between-trial variability: The racing diffusion model of speeded decision making. *Psychonomic Bulletin & Review*, 27:911–936, 2020.

[400] Gasper Tkacik, Elad Schneidman, II Berry, J Michael, and William Bialek. Ising models for networks of real neurons. *arXiv preprint q-bio/0611072*, 2006.

[401] Philippe N Tobler, John P O'Doherty, Raymond J Dolan, and Wolfram Schultz. Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *Journal of neurophysiology*, 97(2):1621–1632, 2007.

[402] Alessandro Tomassini, Darren Price, Jiaxiang Zhang, and James B Rowe. On the evolution of neural decisions from uncertain visual input to uncertain actions. *bioRxiv*, page 803049, 2019.

[403] Giulio Tononi. Consciousness as integrated information: a provisional manifesto. *The Biological Bulletin*, 215(3):216–242, 2008.

[404] Jonathan Touboul, Fabrice Wendling, Patrick Chauvel, and Olivier Faugeras. Neural mass activity, bifurcations, and epilepsy. *Neural computation*, 23(12):3232–3286, 2011.

[405] Dorothée G. A. Kasteleijn-Nolst Trenité, Bettina Schmitz, . . . , Kazuhiro Yamakawa, and Pierre Genton. Consensus on diagnosis and management of JME: From founder's observations to current trends. *Epilepsy & Behavior*, 28:S87 – S90, 2013. ISSN 1525-5050. doi: https://doi.org/10.1016/j.yebeh.2012.11.051. URL http://www.sciencedirect.com/science/article/pii/S1525505012007457.

[406] Stefan Treue and Julio C Martinez Trujillo. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, 399(6736):575–579, 1999.

[407] Tripp, Bryan. A search for principles of basal ganglia function, 2009. URL http://hdl.handle.net/10012/4179.

[408] Jennifer S Trueblood, Scott D Brown, and Andrew Heathcote. The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological review*, 121(2):179, 2014.

[409] Amos Tversky and Daniel Kahneman. Rational choice and the framing of decisions. In *Multiple criteria decision making and risk analysis using microcomputers*, pages 81–126. Springer, 1989.

[410] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323, 1992.

[411] M. Twain and E.P. Hingston. *The Innocents Abroad*. Number vol. 2 in Library of Classics. Collins, 1869. URL https://books.google.ie/books?id=8-RTAAAAcAAJ.

[412] Deirdre M Twomey, Peter R Murphy, Simon P Kelly, and Redmond G O'Connell. The classic p300 encodes a build-to-threshold decision variable. *European journal of neuroscience*, 42(1):1636–1643, 2015.

[413] N. Tzourio-Mazoyer, B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, and M. Joliot. Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain. *NeuroImage*, 15(1):273 – 289, 2002. ISSN 1053-8119. doi: https://doi.org/10.1006/nimg.2001.0978. URL http://www.sciencedirect.com/science/article/pii/S1053811901909784.

[414] Marius Usher and James L. McClelland. The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3):550–592, 2001.

[415] Rens van de Schoot, Sarah Depaoli, Ruth King, Bianca Kramer, Kaspar Märtens, Mahlet G Tadesse, Marina Vannucci, Andrew Gelman, Duco Veen, Joukje Willemsen, et al. Bayesian statistics and modelling. *Nature Reviews Methods Primers*, 1(1): 1–26, 2021.

[416] Tim Van Gelder. What might cognition be, if not computation? *The Journal of Philosophy*, 92(7):345–381, 1995.

[417] Peter van Inwagen. When is the will free? *Philosophical Perspectives*, 3:399–422, 1989. ISSN 15208583, 17582245. URL http://www.jstor.org/stable/2214275.

[418] Marieke Karlijn van Vugt, Patrick Simen, Leigh E Nystrom, Philip Holmes, and Jonathan D Cohen. EEG oscillations reveal neural correlates of evidence accumulation. *Frontiers in neuroscience*, 6:106, 2012.

[419] Joachim Vandekerckhove, Francis Tuerlinckx, and Michael D Lee. Hierarchical diffusion models for two-choice response times. *Psychological methods*, 16(1):44, 2011.

[420] Aki Vehtari, Andrew Gelman, and Jonah Gabry. Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and computing*, 27(5):1413–1432, 2017.

[421] Rolf Verleger, Piotr Jaśkowski, and Edmund Wascher. Evidence for an integrative role of p3b in linking reaction to perception. *Journal of Psychophysiology*, 19(3):165–181, 2005.

[422] Aaron Russell Voelker, Eric Crawford, and Chris Eliasmith. Learning large-scale heteroassociative memories in spiking neurons. *Unconventional Computation and Natural Computation*, 7, 2014.

[423] Nina Vogt. Machine learning in neuroscience. *Nature Methods*, 15(1):33–33, Jan 2018. ISSN 1548-7105. doi: 10.1038/nmeth.4549. URL https://doi.org/10.1038/nmeth.4549.

[424] Katharina Voigt, Carsten Murawski, Sebastian Speer, and Stefan Bode. Hard decisions shape the neural coding of preferences. *Journal of Neuroscience*, 39(4):718–726, 2019.

[425] C. Vollmar, J. O'Muircheartaigh, M.R Symms, G.J. Barker, P. Thompson, V. Kumari, J. Stretton, J.S. Duncan, M.P. Richardson, and M.J. Koepp. Altered microstructural connectivity in juvenile myoclonic epilepsy. *Neurology*, 78(20):1555–1559, 2012. ISSN 0028-3878. doi: 10.1212/WNL.0b013e3182563b44. URL http://n.neurology.org/content/78/20/1555.

[426] Christian Vollmar, Jonathan O'Muircheartaigh, Gareth J. Barker, Mark R. Symms, Pamela Thompson, Veena Kumari, John S. Duncan, Dieter Janz, Mark P. Richardson, and Matthias J. Koepp. Motor system hyperconnectivity in juvenile myoclonic epilepsy: a cognitive functional magnetic resonance imaging study. *Brain*, 134(6):1710–1719, 05 2011. ISSN 0006-8950. doi: 10.1093/brain/awr098. URL https://doi.org/10.1093/brain/awr098.

[427] Jiri Vrba and Stephen E. Robinson. Signal Processing in Magnetoencephalography. *Methods*, 25(2):249 – 271, 2001. ISSN 1046-2023. doi: https://doi.org/10.1006/meth.2001.1238. URL http://www.sciencedirect.com/science/article/pii/S1046202301912381.

[428] Mai-Anh T Vu, Tülay Adalı, Demba Ba, György Buzsáki, David Carlson, Katherine Heller, Conor Liston, Cynthia Rudin, Vikaas S Sohal, Alik S Widge, et al. A shared vision for machine learning in neuroscience. *Journal of Neuroscience*, 38(7):1601–1607, 2018.

[429] Eric-Jan Wagenmakers, Jonathon Love, Maarten Marsman, Tahira Jamil, Alexander Ly, Josine Verhagen, Ravi Selker, Quentin F Gronau, Damian Dropmann, Bruno Boutin, et al. Bayesian inference for psychology. part ii: Example applications with jasp. *Psychonomic bulletin & review*, 25(1):58–76, 2018.

[430] Ben Jonathan Wagner, Mareike Clos, Tobias Sommer, and Jan Peters. Dopaminergic modulation of human inter-temporal choice: a diffusion model analysis using the d2-receptor-antagonist haloperidol. *bioRxiv*, 2020.

[431] A. Wald and J. Wolfowitz. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 19(3):326–339, 1948. ISSN 00034851. URL http://www.jstor.org/stable/2235638.

[432] Abraham Wald and Jacob Wolfowitz. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, pages 326–339, 1948.

[433] Xiao-Jing Wang. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5):955–968, 2020/04/28 2002.

[434] Yanlu Wang and Tie-Qiang Li. Analysis of whole-brain resting-state fmri data using hierarchical clustering approach. *PloS one*, 8(10):e76315, 2013.

[435] Takamitsu Watanabe and Geraint Rees. Brain network dynamics in high-functioning individuals with autism. *Nature Communications*, 8(1):1–14, 2017.

[436] Takamitsu Watanabe, Satoshi Hirose, Hiroyuki Wada, Yoshio Imai, Toru Machida, Ichiro Shirouzu, Seiki Konishi, Yasushi Miyashita, and Naoki Masuda. A pairwise maximum entropy model accurately describes resting-state human brain networks. *Nature Communications*, 4:1370 – 1380, Jan 2013. URL https://doi.org/10.1038/ncomms2388. Article.

[437] Takamitsu Watanabe, Satoshi Hirose, Hiroyuki Wada, Yoshio Imai, Toru Machida, Ichiro Shirouzu, Seiki Konishi, Yasushi Miyashita, and Naoki Masuda. Energy landscapes of resting-state brain networks. *Frontiers in Neuroinformatics*, 8:12, 2014. ISSN 1662-5196. doi: 10.3389/fninf.2014.00012. URL https://www.frontiersin.org/article/10.3389/fninf.2014.00012.

[438] Takamitsu Watanabe, Naoki Masuda, Fukuda Megumi, Ryota Kanai, and Geraint Rees. Energy landscape and dynamics of brain activity during human bistable perception. *Nature Communications*, 5(1):4765, 2014. ISSN 2041-1723. doi: 10.1038/ncomms5765. URL https://doi.org/10.1038/ncomms5765.

[439] Carter Wendelken, Jochen Ditterich, Silvia A Bunge, and Cameron S Carter. Stimulus and response conflict processing during perceptual decision making. *Cognitive, Affective, & Behavioral Neuroscience*, 9(4):434–447, 2009.

[440] David R Wheeler. Brand loyalties: Qualitative, quantitative, or both? *Journal of the Academy of Marketing Science*, 2(4):651–658, 1974.

[441] Corey N White, Eliza Congdon, Jeanette A Mumford, Katherine H Karlsgodt, Fred W Sabb, Nelson B Freimer, Edythe D London, Tyrone D Cannon, Robert M Bilder, and Russell A Poldrack. Decomposing decision components in the stop-signal task: a model-based approach to individual differences in inhibitory control. *Journal of cognitive neuroscience*, 26(8):1601–1614, 2014.

[442] James CR Whittington and Rafal Bogacz. Theories of error back-propagation in the brain. *Trends in Cognitive Sciences*, 23(3):235–250, 2019.

[443] Szymon Wichary and Tomasz Smolen. Neural underpinnings of decision strategy selection: a review and a theoretical model. *Frontiers in neuroscience*, 10:500, 2016.

[444] Szymon Wichary, Rui Mata, and Jörg Rieskamp. Probabilistic inferences under emotional stress: how arousal affects decision processes. *Journal of Behavioral Decision Making*, 29(5), 2016.

[445] Szymon Wichary, Mikołaj Magnuski, Tomasz Oleksy, and Aneta Brzezicka. Neural signatures of rational and heuristic choice strategies: a single trial erp analysis. *Frontiers in human neuroscience*, 11, 2017.

[446] Wayne A Wickelgren. Speed-accuracy tradeoff and information processing dynamics. *Acta psychologica*, 41(1):67–85, 1977.

[447] Thomas Wiecki, Imri Sofer, and Michael Frank. Hddm: Hierarchical bayesian estimation of the drift-diffusion model in python. *Frontiers in Neuroinformatics*, 7:14, 2013. ISSN 1662-5196. doi: 10.3389/fninf.2013.00014. URL https://www.frontiersin.org/article/10.3389/fninf.2013.00014.

[448] Thomas Wiecki, Imri Sofer, and Michael J Frank. Hddm: hierarchical bayesian estimation of the drift-diffusion model in python. *Frontiers in neuroinformatics*, 7:14, 2013.

[449] Hugh R Wilson and Jack D Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12(1):1–24, 1972.

[450] Tomasz Witkowski. *Psychology gone wrong: The dark sides of science and therapy*. Universal Publishers, 2015.

[451] Peter Wolf and Sándor Beniczky. Understanding ictogenesis in generalized epilepsies. *Expert Review of Neurotherapeutics*, 14(7):787–798, 2014. doi: 10.1586/14737175.2014.925803. URL https://doi.org/10.1586/14737175.2014.925803.

[452] Peter Wolf, Elza Márcia Targas Yacubian, Giuliano Avanzini, Thomas Sander, Bettina Schmitz, Britta Wandschneider, and Matthias Koepp. Juvenile myoclonic epilepsy: A system disorder of the brain. *Epilepsy Research*, 114:2 – 12, 2015. ISSN 0920-1211. doi: https://doi.org/10.1016/j.eplepsyres.2015.04.008. URL http://www.sciencedirect.com/science/article/pii/S0920121115000819.

[453] Bo-Yeong Won and Andrew B Leber. How do magnitude and frequency of monetary reward guide visual search? *Attention, Perception, & Psychophysics*, 78(5):1221–1231, 2016.

[454] Kong-Fatt Wong and Xiao-Jing Wang. A recurrent network mechanism of time integration in perceptual decisions. *Journal of Neuroscience*, 26(4):1314–1328, 2006. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.3733-05.2006. URL https://www.jneurosci.org/content/26/4/1314.

[455] Kong-Fatt Wong, Alexander C Huk, Michael N Shadlen, and Xiao-Jing Wang. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Frontiers in Computational Neuroscience*, 1:6, 2007.

[456] KongFatt Wong-Lin, Jose M. Sanchez-Bornot, N. McCombe, Daman Kaur, P. McClean, Xin Zou, V. Youssofzadeh, Xuemei Ding, M. Bucholc, Su Yang, G. Prasad, D. Coyle, Liam P. Maguire, Haiying Wang, H. Wang, Nadim A. A. Atiya, and Alok Joshi. Computational neurology: Computational modeling approaches in dementia. *arXiv: Neurons and Cognition*, pages 81–89, 2020.

[457] Valentin Wyart, Nicholas E. Myers, and Christopher Summerfield. Neural mechanisms of human perceptual choice under focused and divided attention. *Journal of neuroscience*, 35(8):3485–3498, Feb 2015. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.3276-14.2015. URL https://doi.org/10.1523/JNEUROSCI.3276-14.2015.

[458] Elza Márcia Yacubian and Peter Wolf. Praxis induction. definition, relation to epilepsy syndromes, nosological and prognostic significance. a focused review. *Seizure*, 23(4):247 – 251, 2014. ISSN 1059-1311. doi: https://doi.org/10.1016/j.seizure.2014.01.011. URL http://www.sciencedirect.com/science/article/pii/S1059131114000247.

[459] Motonori Yamaguchi, Matthew JC Crump, and Gordon D Logan. Speed-accuracy trade-off in skilled typewriting: Decomposing the contributions of hierarchical control loops. *Journal of Experimental Psychology: Human Perception and Performance*, 39 (3):678, 2013.

[460] LC Yan, B Yoshua, and H Geoffrey. Deep learning. *Nature*, 521(7553):436–444, 2015.

[461] Xin-hua Yang, Jia Huang, Yong Lan, Cui-ying Zhu, Xiao-qun Liu, Ye-fei Wang, Eric FC Cheung, Guang-rong Xie, and Raymond CK Chan. Diminished caudate and superior temporal gyrus responses to effort-based decision making in patients with first-episode major depressive disorder. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, 64:52–59, 2016.

[462] Fang-Chin Yeh, Aonan Tang, Jon P. Hobbs, Pawel Hottowy, Wladyslaw Dabrowski, Alexander Sher, Alan Litke, and John M. Beggs. Maximum entropy approaches to living neural networks. *Entropy*, 12(1):89–106, 2010. ISSN 1099-4300. doi: 10.3390/e12010089. URL http://www.mdpi.com/1099-4300/12/1/89.

[463] Michael E Young, Tara L Webb, Jillian M Rung, and Anthony W McCoy. Outcome probability versus magnitude: When waiting benefits one at the cost of the other. *PloS one*, 9(6), 2014.

[464] Shan Yu, Hongdian Yang, Hiroyuki Nakahara, Gustavo S. Santos, Danko Nikolić, and Dietmar Plenz. Higher-order interactions characterized in cortical activity. *Journal of Neuroscience*, 31(48):17514–17526, 2011. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.3127-11.2011. URL https://www.jneurosci.org/content/31/48/17514.

[465] Wojciech Zajkowski, Dominik Krzemiński, Jacopo Barone, Lisa H. Evans, and Jiaxiang Zhang. Breaking deadlocks: Reward probability and spontaneous preference shape voluntary decisions and electrophysiological signals in humans. *Computational Brain & Behavior*, Nov 2020. ISSN 2522-087X. doi: 10.1007/s42113-020-00096-6. URL https://doi.org/10.1007/s42113-020-00096-6.

[466] Laura Zamarian, Julia Höfler, Giorgi Kuchukhidze, Margarete Delazer, Elisabeth Bonatti, Georg Kemmler, and Eugen Trinka. Decision making in juvenile myoclonic epilepsy. *Journal of Neurology*, 260(3):839–846, Mar 2013. ISSN 1432-1459. doi: 10.1007/s00415-012-6715-z. URL https://doi.org/10.1007/s00415-012-6715-z.

[467] Huaxiang Zhang and Zhichao Wang. A normal distribution-based over-sampling approach to imbalanced data classification. In *International Conference on Advanced Data Mining and Applications*, pages 83–96. Springer, 2011.

[468] Jiaxiang Zhang. The effects of evidence bounds on decision-making: theoretical and empirical developments. *Frontiers in psychology*, 3:263, 2012.

[469] Jiaxiang Zhang and Rafal Bogacz. Optimal decision making on the basis of evidence represented in spike trains. *Neural Computation*, 22:1113–1148, 2010. ISSN 08997667. doi: 10.1162/neco.2009.05-09-1025.

[470] Jiaxiang Zhang and James B. Rowe. Dissociable mechanisms of speed-accuracy tradeoff during visual perceptual learning are revealed by a hierarchical drift-diffusion model. *Frontiers in Neuroscience*, 8:69, 2014. ISSN 1662-453X. doi: 10.3389/fnins.2014.00069. URL https://www.frontiersin.org/article/10.3389/fnins.2014.00069.

[471] Jiaxiang Zhang and James B Rowe. The neural signature of information regularity in temporally extended event sequences. *Neuroimage*, 107:266–276, 2015.

[472] Jiaxiang Zhang, Timothy Rittman, Cristina Nombela, Alessandro Fois, Ian Coyle-Gilchrist, Roger A Barker, Laura E Hughes, and James B Rowe. Different decision deficits impair response inhibition in progressive supranuclear palsy and parkinson's disease. *Brain*, 139(1):161–173, 2016.

[473] Lin Zhang, Hua Li, Peiwei Hong, and Xiaoyi Zou. Proton magnetic resonance spectroscopy in juvenile myoclonic epilepsy: A systematic review and meta-analysis. *Epilepsy Research*, 121:33 – 38, 2016. ISSN 0920-1211. doi: https://doi.org/10.1016/j.eplepsyres.2016.01.004. URL http://www.sciencedirect.com/science/article/pii/S0920121116300043.

[474] Xiang Zhou, KongFatt Wong-Lin, and Holmes Philip. Time-varying perturbations can distinguish among integrate-to-threshold models for perceptual decision making in reaction time tasks. *Neural computation*, 21(8):2336–2362, 2009.

[475] Chenzhuo Zhu, Song Han, Huizi Mao, and William J Dally. Trained ternary quantization. *arXiv preprint arXiv:1612.01064*, 2016.