

Seeing the forest through the
trees: Structure representation
under noisy processing

Fabrice Luyckx

New College

Department of Experimental Psychology

University of Oxford

United Kingdom

A thesis presented for the degree of
Doctor of Philosophy in Experimental Psychology

Hilary term 2020

Abstract

For humans to optimally handle the constant stream of input they receive from their surroundings, they need to be able to separate signal from noise. The understanding that information is structured can help humans to organize information better and improve the detection of signal in a noisy input. This thesis deals with two main themes: the neural representation of structure and how structure representation can adapt in the light of noisy processing. First, I investigate how the human brain represents the structure of magnitude and how through structural alignment it can learn efficiently about new stimuli that have the same underlying structure. Next I present evidence from a new task that magnitude is represented neurally along a single axis but separated in parallel lines based on the current context. In the following chapter I investigate through model simulations how normalization can confer robustness under accumulation noise and test these predictions in humans. Finally, I present neural evidence for a computational model of robust processing when two pieces of information have to be assessed simultaneously. Taken together I hope to demonstrate that neural signals carry a low-dimensional representation of structure and that the brain can adaptively change those representations to improve performance.

Acknowledgements

This doctorate would have never made it to completion without the help, advice, love and understanding of a bunch of people. I will most certainly forget to mention a lot of them and to them I apologize in advance.

First and foremost, I would like to thank **prof. Christopher Summerfield** for all he has done for me throughout the last five years. I have come a long way from being that naive Erasmus student to the scientist I am today, and I owe most (if not all) to you. You are truly a supervisor that any graduate student can only dream of. You always have the best interest of your students in mind and try to elevate them to higher levels. Also, keep rocking those awesome sweaters.

The Summerfield lab with all its current and former members have been some of the most wonderful colleagues I could have imagined. And while I will single out some, each and every one of you has a special place in my heart. **Keno**, some of my fondest memories during my time in the lab have been with you. From the start you were always there when I needed advice. We have travelled to some amazing places and I am glad to have shared those experiences with you. Thank you for being such a great colleague, scientist but most of all, friend. **Timo**, our paths have run parallel for most of our time in Oxford. From the early days in Christian housing to the slightly fancier New College, you have always been a friend that I could count on

if in need of a drink (or technical questions). I cannot think of a better metaphor to our friendship than *Jägerbombs*: full of caffeine and alcohol (and the occasional science of course). Cheers! **Hannah T**, you sweet potato (or something along those lines). Even though you were not around that often, you have certainly made your mark on my time in Oxford. You were always there when I needed advice, or gossip. You are a strong, fierce, funny and independent woman and for that I will always admire you.

I would also like to thank **prof. Tom Verguts** and his lab at Ghent University to give me shelter in those dark days of thesis writing. You have been extremely accommodating and I am sure the struggle of writing would have been a lot more difficult without such a stimulating work environment.

Sam and **Joshua**, my brothers in arms. I am sad to think that our paths will take different directions after this doctorate, but our time together getting through this doctorate has been amazing and I am glad that I was able to share the highs and the lows with you. You are brilliant scientists with a bright future ahead. My friends from Marston Manor and Regent Palace, **Victoire**, **Mirjam**, **Aaron**, **Naveen** and (again) **Sam**, living with you has made my days in Oxford so much brighter. One day I will kick your asses at Catan again.

Sofie, I cannot begin to imagine how I would have made it through without your support day and night. You have had to sit through all the good and the bad parts of this doctorate. Thank you for sticking by my side the entire time and making my life so much more pleasant.

Mama, **papa**, you have had to watch your son leave for England and sit at the sideline while I was doing my thing. But I would have never made it this far without your unconditional love and support. You are the best.

Statement of Contribution

The quality of the work presented in this thesis owes much to the contributions of collaborators throughout the various projects. All empirical chapters were developed under the supervision of prof. Christopher Summerfield. In particular the neural network simulations reported in Chapter 2 (p. 59) were developed by prof. Summerfield. The line normalization model from Chapter 3 (p. 83) was similarly initiated by prof. Summerfield during our work on the project. Dr. Bernhard Spitzer's contribution to the EEG side of Chapters 2 and 5 has been extremely valuable. The multivariate analyses (p. 43) and psychometric model (p. 46) reported in Chapter 2 were based on original code by Spitzer et al. (2017), but were adapted by myself to suit the purposes of the project. For completeness some projects would not have been what they are today without the help in some form by the following people: Hannah Sheahan (Chapters 1,6), Hamed Nili (Chapter 2), Clemens Teupe (Chapter 3), Annabelle Blangero (Chapter 5) and Konstantinos Tsetsos (Chapter 5).

The following chapter have been published in scientific journals: Chapter 2 (<https://doi.org/10.7554/eLife.42816.001>) and Chapter 5 (<https://doi.org/10.1093/cercor/bhaa039>). All co-authors have agreed to include the material in the thesis.

Contents

Abstract	iii
Acknowledgements	v
Statement of Contribution	vii
1 General Introduction	1
1.1 What is structure?	3
1.1.1 Early views on structure representation	4
1.1.2 Characteristics of structure	6
1.2 Models of structure learning	8
1.2.1 Probabilistic models	9
1.2.2 Connectionist models	11
1.2.3 Reinforcement learning models	13
1.3 Extracting representational structure from neural signals . . .	16
1.4 Robust processing of information	18
1.5 Concluding remarks	20
References	22
2 Neural structure mapping in human probabilistic reward learning	33

2.1	Introduction	33
2.2	Methods	35
2.2.1	Participants	35
2.2.2	Experimental procedure	36
2.2.3	EEG acquisition	40
2.2.4	EEG pre-processing	40
2.2.5	Statistical procedure	41
2.2.6	Delta-rule model	42
2.2.7	Representational Similarity Analysis (RSA)	43
2.2.8	CPP analysis	45
2.2.9	Dimensionality reduction	45
2.2.10	Psychometric model	46
2.2.11	Neurometric mapping	47
2.2.12	Neural network	48
2.3	Results	49
2.3.1	Shared magnitude representation for numbers and probabilistic rewards	49
2.3.2	Relating numerical magnitude representation to choice behaviour	53
2.3.3	Dimensionality of magnitude representation	57
2.3.4	Neural network simulations	59
2.4	Discussion	62
	References	65
3	Aligned neural geometry between numerical contexts in humans	71
3.1	Introduction	71
3.2	Methods	73

3.2.1	Participants	73
3.2.2	Experimental procedure	74
3.2.3	Behavioural analyses	77
3.2.4	EEG acquisition	78
3.2.5	EEG pre-processing	78
3.2.6	EEG: Univariate analyses	79
3.2.7	EEG: Representational Similarity Analysis (RSA) . . .	80
3.2.8	Line configuration model	81
3.2.9	Line normalization model	83
3.2.10	Cross-generalisation analysis	84
3.3	Results	85
3.3.1	Context-driven processing	85
3.3.2	Multivariate representation of context and magnitude	88
3.3.3	Aligned geometry in neural signals	90
3.4	Discussion	94
	References	96
4	Robust decision-making through normalization and compression under noisy processing	99
4.1	Introduction	99
4.2	Methods	102
4.2.1	Participants	102
4.2.2	Experimental procedure	102
4.2.3	Psychometric models and model fitting	104
4.2.4	Model simulations	105
4.3	Results	107
4.3.1	Model simulations	107
4.3.2	Interim discussion	110

4.3.3	Model fits to human data	110
4.4	Discussion	112
	References	114
5	Selective integration during sequential sampling in posterior neural signals	117
5.1	Introduction	117
5.2	Methods	120
5.2.1	Participants	120
5.2.2	Experimental procedure	120
5.2.3	Selective Integration model	125
5.2.4	Model fitting procedure	127
5.2.5	Behavioural analyses: regression models	128
5.2.6	EEG acquisition	129
5.2.7	EEG pre-processing	130
5.2.8	Time-frequency transformation	131
5.2.9	EEG analyses: regression models	131
5.3	Results	136
5.3.1	Behavioural results	136
5.3.2	Model fits to human data	138
5.3.3	Early modulation of posterior EEG activity	139
5.3.4	Alpha suppression	143
5.3.5	Evidence accumulation	144
5.4	Discussion	145
	References	149
6	General Discussion	153
6.1	Structure learning models revisited	153

6.2	Posterior parietal cortex as candidate region	157
6.3	Flexible adaptive structure representation?	162
6.4	Concluding remarks	164
	References	166
A	Supplementary material to Chapter 2	171
B	Supplementary material to Chapter 3	181
C	Supplementary material to Chapter 4	185
D	Supplementary material to Chapter 5	189

List of Figures

1.1	Structure in noise	2
1.2	Gestalt examples	4
1.3	Structure from content	7
1.4	Structure examples	8
2.1	Task design and RSA results	37
2.2	Cross-temporal cross-validation RSA	50
2.3	Detailed examination of the magnitude pattern in cross-validation	52
2.4	Behavioural analyses	54
2.5	Dimensionality of magnitude representation	58
2.6	Neural network simulations	60
3.1	Task design	75
3.2	Context-dependent processing	87
3.3	Magnitude and context in the brain	89
3.4	Aligned geometry	92
4.1	Task design	103
4.2	Accuracy under late noise with fixed k	108
4.3	Accuracy under late noise with multiple k	109
4.4	Model fitting results	111

5.1	Task design and model schematic	123
5.2	Behavioural results and model fits	135
5.3	Early modulation by relative sample evidence in posterior electrodes	140
5.4	Encoding of evidence integration in central electrodes	144
6.1	Examples of scalar codes	159
6.2	Low-dimensional coding in human PPC	161
A.1	RSA with correlation distance	172
A.2	RSA results split per task framing	173
A.3	RSA results split for first and second bandit	174
A.4	Cross-validation after removal of n6 - b6	175
A.5	Testing strength of within-subject cross-validation	176
A.6	CPP by task and task framing	177
A.7	Controlling for CPP in multivariate results	178
A.8	Cross-validation after excluding the first principal dimension .	179
B.1	RSA control analysis	182
B.2	Examples configuration model	183
C.1	Model fits for mixed trials	186
C.2	Assessing inflection point around decades	187
D.1	Parameter recovery	190
D.2	HEOG control analysis	192
D.3	EEG control analysis	194

Chapter 1

General Introduction

In 1965, photographer R. C. James published a photograph in Life Magazine that half a century later would still be used in every introductory Psychology class (Figure 1.1)¹. To the naive observer the picture initially consists of black and white patches scattered across the frame. But what if they are then told a dog hides in the picture? Not only will the contours of a dog start to emerge, it will be hard to ever see the image again without noticing the dog.

This phenomenon exemplifies a fundamental conundrum for humans in everyday life: How do we make sense of our surroundings? The problem may be best described by psychology's great-grandfather William James: *"Certainly a child newly born in Boston, who gets a sensation from the candle-flame which lights the bedroom, or from his diaper-pin, does not feel either of these objects to be situated in longitude 72° W. and latitude 41° N. [...] He does not, in short, know anything about their space-relations to anything else in the world. The flame fills its own place, the pain fills its own place; but as yet these places are neither identified with, nor discriminated*

¹James, R.C. (1965), *Sight for sharp eyes*. Life Magazine (1965-02-15), p. 120



Figure 1.1. Structure in noise. One of the fundamental problems humans face every day is to distinguish signal from noise. Initially the image might appear to contain only noisy black and white patches. But if we are instructed to expect signal, we can impose known structural regularities onto the image and it becomes hard to ever notice the Dalmatian again. Our ability to use our knowledge of structure can help us distinguish signal from noise.

from, any other places." (James, 1890, p. 681-682). To James' child, the world is like a picture of black and white patches with no order or meaning. It is only later that it will learn the world is full of structure, allowing for meaningful interpretation of all those noisy signals.

As we mature, we start to naturally seek out structure in our environment, even when it sometimes leads to illusory perceptions. Staring at the clouds, one often distinguishes faces, ships or other shapes (called 'pareidolia'). The same phenomenon has taught us that Jesus Christ prefers to re-appear to the material world on burnt toast (Liu et al., 2014). But more often the tendency to assume structure in the world is a useful trait and an important aspect of human intelligence (Gentner, 2003, 2010). It allows us to make inferences that go beyond what can be observed directly (Alfred, Connolly, & Kraemer, 2018; Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010;

Kemp & Tenenbaum, 2009). When all the bins in your street are outside, you can assume that bin collection is due the next day, because those events are generally linked. We can also make predictions about the future or what might happen in new situations. When we get sick after eating some new delicacy while travelling, we can infer that we should probably not eat that food again. This in turn can save useful energy resources; for example we do not have to worry each night whether the sun will come up the next day, because we understand there is structure in the diurnal cycle.

If we as psychologists want to understand human behaviour better, a good starting point thus seems to be the study of how humans represent structure and the role of structure representation in information-processing. Over the course of four empirical chapters (2 - 5), I have taken a neural approach (EEG) in combination with computational modelling to answer these questions. In the following paragraphs, I will attempt to define the concept of 'structure', introduce several models of structure learning and their contribution to understanding the underlying neural mechanisms and finally briefly review what we know about adaptive information-processing given a noisy neural system.

1.1 What is structure?

To study structure representation, we first need to delimit what is meant by 'structure'. Throughout history, the conceptualization of structure representation and structure learning has taken many forms. I discuss some of these influential historical ideas and try to bring them together to obtain a working definition of 'structure'.




Figure 1.2. Examples of Gestalt configurations. The principle of *proximity* states that elements are grouped when they are closer together. With *similarity* elements resembling each other in some way, such as size or colour, will be grouped, independent of other features. When elements that are organized according to a regular pattern, such as a line, are perceived as coming from the same object, they follow the principle of *alignment*. Grouping by *closure* occurs when elements form a closed figure (adapted from Summerfield et al. (2019)).

1.1.1 Early views on structure representation

Nearly a century ago, Max Wertheimer and the Gestalt movement understood that perception is not objective, but coloured by our expectations and prior knowledge (Wagemans et al., 2012). They defined several principles under which elements are perceived as parts of the same object, so-called 'principles of grouping' (see Figure 1.2 for several examples). For principles such as *alignment* or *closure*, this involves identifying regular patterns and inferring that items following these patterns must come from the same object. Crucially grouping principles can take priority over other shared visual features, such as size or colour. One could even replace the elements entirely, while the perceptual experience of grouped items would remain constant.

How do we come to perceive structure in the world? Jean Piaget, a developmental psychologist in Geneva, Switzerland, proposed an answer to William James' problem by introducing the concept of 'schemata' or 'cognitive

structures'. According to Piaget schemata are the building blocks to organize knowledge and to build a cognitive model of the world: "... *by constructing [...] schemata such as those of the unchanging object, the fitting in of two or three dimensions, rotations, transpositions, and superpositions that he finally succeeds in organizing his 'mental space', ...*" (Piaget, 1952, p. x). Initially, a schema for 'dog' can consist of features such as 'has four legs' and 'has a tail'. Over time schemata can become more complex and can change when faced with new situations. Upon encountering a cat for the first time, the child might mislabel it as 'dog', but can learn new features that distinguish cats from dogs.

In the hands of Edward Tolman (1948), structured knowledge was called a 'cognitive map'. In Tolman's original experiments rats would learn to follow a certain path in a maze to obtain a reward. After learning the rats were placed in a new maze that had more direct routes to the same reward. If rats just associated direction with reward, they would continue to follow the trajectory learned in the old maze. However, if they had developed an understanding of their environment that was independent of their original trajectory, they could apply that knowledge to the new configuration and choose a more direct route. This is indeed what Tolman found. The implication here is that cognitive maps allow for flexible behaviour (Behrens et al., 2018), or as Tolman put it: "... *the wider and the more comprehensive it [the map] was, the more adequately it will serve in the new set-up.*" (Tolman, 1948, p. 193). And although the original experiments dealt with spatial navigation, Tolman in his discussion already envisaged a broader application of cognitive maps, for example in social settings.

Along similar lines to Piaget and Tolman, Dedre Gentner's theory of 'structure-mapping' (1983) incorporated the notion that concepts derive their

meaning from their relation to other concepts. The core idea is that children learn through analogical reasoning. An analogy is a system of relations that maps one set of elements ('base') onto a different set ('target'). The base and target sets do not need to resemble each other in any way, as long as they follow the same relational structure. For example, children can be taught in school that atoms work in a similar fashion to our solar system: electrons orbit the nucleus like planets orbit the Sun. In order to understand the analogy, a child must first be able to abstract the structure (X orbits around Y) from the elements themselves (Earth and Sun), before it can apply that structure to a new set of elements (electrons and nucleus).

1.1.2 Characteristics of structure

All the above views point towards an important common feature: structure is orthogonal to the elements to which it applies (Figure 1.3). In other words, we need to make a clear distinction between elements on the one hand and the associations among them on the other hand. Elements can be anything from tangible physical items such as cars parked on the side of the road, to more abstract concepts like countries in Europe. Similarly structures can go from simple low-dimensional forms to more complex or abstract forms. Numbers or the alphabet can be placed on a line (1-dimensional), weekdays or seasons follow a circular pattern and the geolocations of cities on Earth are positioned on a three-dimensional sphere. More abstractly, family trees or companies are organised hierarchically and even more complex are the rules that describe how to play a game of chess or the recipe to making the perfect scrambled eggs (Figure 1.4). These examples should make it clear that associations take many forms: spatial ('A is closer to B than to C'), temporal ('After A, B happens'), causal ('When A is pushed, B lights up'),




Figure 1.3. A visual scene is understood not just by its contents but by the relations between the elements that are present. In the example on the left, scene understanding depends on the relative size of the two elements (man $>$ car or car $>$ man) and containment (man in car or man outside of car) Right panel: a different example with two elements (cats and bowls) and two different structural principles: alignment and cardinality. (Adapted from Summerfield et al. (2019)).

statistical ('Every time A lights up, B lights up too') or other. Of note, in the material world relations between elements can be defined exactly by shared physical properties (see Gestalt), deeming the independence between structure and elements less clear-cut.

If structure is orthogonal to sensory information, it follows that structure representations should be transferable and generalizable to new environments. In line with Tolman and Gentner, several authors have argued that the explicit distinction between sensory information and structure is a necessary requirement for any neural or artificial system to exhibit flexible behaviour (Behrens et al., 2018; Hassabis, Kumaran, Summerfield, & Botvinick, 2017; Lake, Ullman, Tenenbaum, & Gershman, 2017; Summerfield, Luyckx, &




Figure 1.4. Examples of structures. **A.** Numbers or letters in the alphabet lie on a one-dimensional line. **B.** Seasons follow a circular pattern (adapted from “season” by Adrien Coquet, from the Noun Project). **C.** Cities on Earth are located on a three-dimensional sphere (adapted from “Globe” by Designs by MB, from the Noun Project). **D.** Organizations are usually structured hierarchically, with a team leader at the top and several levels of supervision (“Hierarchy” by Adiba Taj, from the Noun Project). **E.** Recipes consist of multiple steps that follow a particular order (adapted from icons by Llisole, from The Noun Project).

Sheahan, 2020). In Chapters 2 and 3, I will present evidence that the human brain does in fact represent structure independent of sensory information or task and generalizes to new elements.

1.2 Models of structure learning

While this thesis focuses on the neural principles of structure learning, most work has come from a computational modelling perspective. In the following

section I will cover different approaches to structure learning and evaluate their contribution to the understanding of the neural mechanisms of structure learning. Traditionally two types of models are pitted against each other on how structure is learned: probabilistic and connectionist models. I will first discuss the two views and then introduce a third approach that in recent years has become relevant for structure learning: reinforcement learning.²

1.2.1 Probabilistic models

Probabilistic models of cognition (Griffiths et al., 2010; Kemp & Tenenbaum, 2008, 2009; Kemp, Tenenbaum, Niyogi, & Griffiths, 2010; Lake et al., 2017; Tenenbaum, Kemp, Griffiths, & Goodman, 2011) adopt a Bayesian approach to structure learning. These models assume a top-down approach where the model starts with a hypothesis space over possible structural forms. By combining observed data with prior beliefs about which forms are more likely, the model can calculate which structures are more probable given the data. The probabilistic nature of these models allows uncertainty over which form is more appropriate and can adapt when confronted with new data. Through a graph generating system, Hierarchical Bayesian models can for example successfully retrieve the tree structure of zoological taxonomy, the political spectrum of judges on the U.S. Supreme Court or geographical locations of world cities on the globe (Kemp & Tenenbaum, 2008). The model also predicts how children over time change from representing words in mutually exclusive clusters to more complex structures where clusters can be hierarchically organised (Kemp & Tenenbaum, 2008). Because learning

²As a side note, the discussion more broadly mirrors the different views in cognitive science. The *symbolic* approach treats the brain as a logical computer. *Connectionist* approaches work on the premise of undifferentiated self-organizing connected units. Finally, the *conceptual* view sees cognition as navigation through a multidimensional conceptual space (Gärdenfors, 2004).

follows Bayes' optimal decision rule, the solutions the model provides can be seen as normative, describing what people *should* do given the observed data (Griffiths et al., 2010; although see Moran and Tsetsos (2018) for a critique on this view).

Despite recent calls that neuroscience should be inspired by state-of-the-art work in computational modelling and machine learning (Hassabis et al., 2017; Richards et al., 2019), remarkably little neural evidence has sprung from these ideas. One of the main criticism of probabilistic models is that it is hard to make mechanistic predictions from their models. From a methodological Bayesian view this is not necessarily problematic, since human behaviour can *appear* to follow Bayes' rule, while the underlying neural system can have very different mechanics that produce a Bayesian-like output (Bowers & Davis, 2012). Nevertheless, numerable theories have been proposed on how Bayesian theory could be implemented neurally (Knill & Pouget, 2004; Ma, Beck, & Pouget, 2008; O'Reilly, Jbabdi, & Behrens, 2012; Tervo, Tenenbaum, & Gershman, 2016; Vilares & Kording, 2011). For example the Bayesian coding hypothesis states that 'the brain represents sensory information probabilistically' (Knill & Pouget, 2004). However probabilistic representation is only the minimal necessary requirement for a brain to be called 'Bayesian'. To be fully Bayesian it would need to follow Bayes' rule: neuronal populations would need to represent the full joint probability space and perform complex integrations over these distributions (Bowers & Davis, 2012; O'Reilly et al., 2012; Rahnev, 2019), which many believe to be highly unlikely (Bowers & Davis, 2012; Moran & Tsetsos, 2018). Despite the theoretical work, empirical evidence is scarce. So far most evidence has come from studies on the visual system that test the proposition that neurons show probabilistic coding (van Bergen, Ma, & Pratte, 2015; Walker,

Cotton, Ma, & Tolias, 2019). To date no studies have directly tested whether Hierarchical Bayesian models of structure learning are implemented in the brain. Moreover, structure learning likely requires the involvement of other functional areas, such as associative areas, that to date have not been studied in the light of these models (Summerfield et al., 2020).

1.2.2 Connectionist models

Often presented as the opposing view to probabilistic models is the connectionist or emergentist view, assuming a bottom-up approach (Doumas, Hummel, & Sandhofer, 2008; McClelland et al., 2010; McClelland & Rogers, 2003; Rumelhart & Todd, 1993; Saxe, McClelland, & Ganguli, 2013, 2019). These models make no *a priori* assumptions about the forms that constitute the final structure. Rather the models mimic a simplified version of how neurons communicate in the brain, consisting of several layers of undifferentiated units, where activity spreads from an initial input layer through one ('shallow') or more ('deep') hidden layers to an output layer. How activity propagates through the network is governed by the decision weights between units in different layers. Decision weights can then be updated through backpropagation after matching the output of the model to the true output. How structure 'emerges' from the system can be studied by analysing the activity patterns in the network.

One influential class of connectionist models is the Parallel Distributed Processing (PDP) approach (Rogers & McClelland, 2004). In the classic case of structure learning the model has two input clusters, one for items (animals, plants, etc.) and one for relation propositions ('can', 'has', etc.). The model gradually learns to make valid completions, i.e. 'salmon can' ... 'swim'. Over time the model starts to cluster items together that show

greater similarity, and does so by first grouping over broad categories and only later starts differentiating within categories. Similar to human children, PDP models can exhibit 'illusory correlations', making inferences on the properties of new items, even when those inferences are not true (Rogers & McClelland, 2004). Another feature that sets connectionist models apart from probabilistic models is so-called 'cross-cutting' (McClelland, Sadeghi, & Saxe, 2016) (although see Shafto, Kemp, Mansinghka, and Tenenbaum, 2011). While probabilistic models will identify the most prominent structure, connectionist models allow items to be part of multiple categories. For example, families can be organized in a tree, but across generations there are other distinctions based on gender or eye colour.

Connectionist models might be more adequate to make predictions about neural mechanisms through their resemblance with real brain networks. On the other hand the opaqueness of neural network learning has often been used as an argument against connectionist approaches, although recently advances have been made in studying the learning dynamics in neural networks (Flesch, Balaguer, Dekker, Nili, & Summerfield, 2018; Hassabis et al., 2017; Saxe et al., 2013, 2019). Several studies have now demonstrated a strong correspondence between early layers in deep convolutional neural networks (DCNN) and early layers of primate primary visual and inferotemporal (IT) cortex (Güçlü & van Gerven, 2015; Khaligh-Razavi & Kriegeskorte, 2014; Kriegeskorte, 2015; Yamins et al., 2014). For example, Khaligh-Razavi and Kriegeskorte (2014) trained several DCNN on images that had also been viewed by humans and monkeys. They then compared the representational structure that emerged in the neural networks to activity in human and monkey IT cortex. DCNN models that most resembled IT representational patterns, i.e. that best captured the categorical nature of the images (animate/inanimate, human/non-human

etc.), also performed better on image classification, suggesting a possible link between the computational features of IT cortex and DCNN.

However, there is no relational information embedded in the connectionist models tested here. Summerfield et al. (2020) suggest the problem is that most DNNs are trained to categorise images based on visual input and are therefore not put under any pressure to learn about structure in the environment. Even though the networks are able to generalize to unseen items, they would likely perform badly when required to perform a new task. Similar to Bayesian models, it seems beyond the visual system little is known about the link between connectionist models and the brain (but see Holroyd, Ribas-Fernandes, Shahnazian, Silvetti, and Verguts (2018) for work on cingulate cortex). Neural network models that were built to explicitly incorporate structure, such as PDP models, have yet to be linked to neural mechanisms or brain areas (Hassabis et al., 2017).

1.2.3 Reinforcement learning models

A third view has recently come to the foreground that could yield promising insights into the learning of structure. At an abstract level cognition is treated as a conceptual space, much like Tolman’s cognitive maps (Behrens et al., 2018; Gärdenfors, 2004). This idea can be formalized in the form of reinforcement learning (RL) (Sutton & Barto, 2018). The goal of a RL agent is to learn how to maximize reward in a certain task. Usually the task is presented as an interactive agent that explores a space of interconnected states (s). Each state can have certain actions (a) associated with it and the transition probabilities determine how states are connected ($P(s'|s, a)$). Relevant for structure learning is that the state representation can be seen as an abstract notion of the underlying structure, namely the understanding

of how states are connected (Behrens et al., 2018; Collins, 2018; Wilson, Takahashi, Schoenbaum, & Niv, 2014).

In model-free RL, the agent gradually learns the value of being in a state through trial-and-error by iteratively visiting states, taking actions and observing what the outcome is. This is a slow process that becomes increasingly intractable as the size of the state space grows ('curse of dimensionality'). Extended versions of RL models have been developed that tackle this problem and further broaden the forms that structure representations can take. Hierarchical RL (HRL) adds a higher-level action space that partitions the state space into subroutines, making large problems more tractable again (Botvinick, 2012; Collins, 2018). Because the action space is abstract, the model can generalize to new items that follow the same task rules, an important feature of structure learning (Collins, 2018). Studies have shown that humans can exploit the hierarchical structure of a task to learn faster compared to a flat solution (Badre, Kayser, & D'Esposito, 2010; Frank & Badre, 2012), and transfer the hierarchical solution to new stimuli (Collins & Frank, 2013).

The close link between physical and conceptual navigation has led to the suggestion that neural systems, originally thought to be responsible only for spatial navigation, could also be used in more abstract navigation (e.g. reasoning problems) (Behrens et al., 2018; Bellmund, Gärdenfors, Moser, & Doeller, 2018; Constantinescu, O'Reilly, & Behrens, 2016; O'Keefe & Nadel, 1980). Work in rodents has identified a whole range of cells in the hippocampal-entorhinal system that seem crucial to navigating physical space. For example, place cells in hippocampus fire only when the animal is at a specific location (O'Keefe & Dostrovsky, 1971), while grid cells in medial entorhinal cortex are location-independent and fire in a hexagonal lattice

across the entire space (Hafting, Fyhn, Molden, Moser, & Moser, 2005). Several studies have now identified signals that suggest a hexagonal grid code also exists in human and non-human primate entorhinal cortex during spatial (Doeller, Barry, & Burgess, 2010; Jacobs et al., 2013; Killian, Jutras, & Buffalo, 2012) and even imagined navigation (Horner, Bisby, Zotow, Bush, & Burgess, 2016). Constantinescu et al. (2016) were the first to identify BOLD signals suggestive of a hexagonal code in human hippocampus and ventromedial prefrontal cortex (vmPFC) in a task where participants had to navigate an abstract 2D feature space defined by the ratio between the neck and leg length of bird drawings. Other studies have since found evidence for navigation through conceptual space in human medial temporal lobe (MTL) (Garvert, Dolan, & Behrens, 2017; Solomon, Lega, Sperling, & Kahana, 2019; Viganò & Piazza, 2020). Solomon et al. (2019) for example found that theta power (4 - 8 Hz) recorded intracranially in MTL varied with the distance between words, not only with their position in the word list but also in semantic space.

In line with hierarchical RL models, mounting evidence now also exists for hierarchical structure representation in the brain (Badre & Frank, 2012; Badre et al., 2010; Collins, Cavanagh, & Frank, 2014; Frank & Badre, 2012; McKenzie et al., 2014). One study found that rodent hippocampus hierarchically organised memories, combining events that shared features while separating events with distinct features (McKenzie et al., 2014). Balaguer, Spiers, Hassabis, and Summerfield (2016) asked participants to navigate a virtual subway network. They found that dorsomedial prefrontal cortex (dmPFC) and anterior premotor cortex (PMC) encoded the distance to goal in terms of subway lines rather than individual stations, an indication that participants used hierarchical planning. Badre and Frank (2012) found

activity was greater in pre-PMC in a condition where hierarchical learning was possible compared to a condition with a flat solution.

In conclusion approaching the problem of structure learning from the perspective of RL has led to a more comprehensive body of findings into the neural mechanisms than with the probabilistic or connectionist models of structure learning. It is also encouraging that structure learning would fit within a larger theoretical framework that encompasses other ideas such as value-based learning, flexible decision-making, memory or rule learning. RL's infrastructure is perhaps also its limiting factor: the computational framework was not primarily intended to serve as an analogue to the neural system. While there is a strong link between RL and neuroscience, for RL models to develop into a more mature model of the neural system, they would need to incorporate for example the noisy processing of neuronal populations, which in turn can alter the way a system learns (for example see (Dabney et al., 2020)).

1.3 Extracting representational structure from neural signals

In recent years the analysis of representational structure in the brain has received a boost with the introduction of a new analysis tool: Representational Similarity Analysis (RSA) (Kriegeskorte & Kievit, 2013; Kriegeskorte, Mur, & Bandettini, 2008; Nili et al., 2014). The basis of RSA (and multivariate techniques in general) is that neurons are treated as different dimensions in a multidimensional space. The same idea can apply to all sorts of 'units', such as neuronal populations, voxels in fMRI, scalp potentials in EEG or other neural measurements. The activity pattern of all units in response to a

certain stimulus can be described as a single point in this multidimensional space. Conventional multivariate techniques such as pattern classifiers utilize the fact that different stimuli or conditions will occupy different positions in the multidimensional space and seek a separation boundary that allows the classification of unseen stimuli. RSA goes one step further and assumes that there is information in the configuration of all points in the representational space. Points that lie closer together are more similar to each other than points that lie further apart. Once all pairwise distances between points ('conditions') are calculated, they are summarized in a Representational Dissimilarity Matrix (RDM). RDMs provide a convenient summary that allows easy comparison to representational patterns in other brain regions, neural networks or model RDMs through simple correlation (Kriegeskorte & Kievit, 2013). Model RDMs correspond to hypotheses about the representational structure contained in the neural signal and can be tested competitively through multiple regression, partial correlation or orthogonalization.

The technique has been applied successfully to anything from single-cell recordings to fMRI or M/EEG (Cichy & Pantazis, 2017). In this thesis RSA is used in combination with EEG. In contrast to fMRI with thousands of voxels, EEG lacks the spatial resolution because of volume conduction and the often limited number of electrodes (usually 64 or 128). Rather than asking *where* a representation might arise, the question becomes *when* representations arise. For example, representations of visual information might occur early, while higher-level or motor-related processes only arise after hundreds of milliseconds.

One study that has greatly influenced the work discussed in this thesis, used RSA in combination with EEG to study the representational geometry of number representation (Spitzer, Waschke, & Summerfield, 2017). Participants

viewed sequences of 10 single-digit numbers between 1 and 6, randomly coloured in green or red. At the end of a sequence participants had to indicate which of the two colour categories contained the highest average. The authors found that around 200 ms after stimulus onset digits 1 to 6 were represented along a line in multivariate neural signals. This was independent of the decision value (a red 6 and green 1 carry equal evidence for responding 'red'), suggesting a more abstract representation of magnitude.

Interestingly, rather than a uniformly spaced number line, larger numbers carried relatively greater weight on choices and the level of non-linear weighting was strongly linked to the distortion found in the neural signals. So far models of structure learning have ignored this aspect of structure representation: they assume that once structure is learned, it can be employed truthfully. However, we know that neuronal firing is noisy and decision-making is a noisy process. This raises the next question that this thesis is concerned with: Can structure representation be changed adaptively?

1.4 Robust processing of information

We know that humans are not good at representing information truthfully. When asked to estimate the distance between cities, humans tend to distort their representation of geographical locations based on geopolitical boundaries, estimating cities from the same county as closer together than those crossing the border (Stevens & Coupe, 1978). The way humans represent quantities is usually not linear either. For example, the same discount of £3 will look very appealing when buying a book of £15, but will hardly raise an eyebrow when buying a new coat at £150. In general, our mental number line appears to be logarithmic instead of linear (Dehaene, 2003), where the subjective difference between two numbers reduces as both numbers increase in size. Humans are

equally 'bad' at estimating probabilities, their judgments are better explained by a non-linear weighting function that overestimates smaller probabilities and underestimates larger probabilities (Kahneman & Tversky, 1979; Prelec, 1998).

Although at first sight these biases might appear to be irrational or suboptimal, instead they might point towards something more fundamental: an efficiency principle (Summerfield & Li, 2019; Summerfield & Tsetsos, 2015). There is a lot of redundant information in our environment. For a system with limited resources to perform efficiently, it makes sense to allocate those resources to the most informative cues in the environment (Barlow, 1961). Originally the theory of efficient coding was proposed to explain adaptation in sensory perception, but increasing evidence suggests the same mechanism could also be used to explain higher-level processes such as decision-making (Li, Michael, Balaguer, Herce Castañón, & Summerfield, 2018; Louie, Khaw, & Glimcher, 2013; Polania, Woodford, & Ruff, 2019; Summerfield & Tsetsos, 2012).

One example is the case of economic decision-making. Research in economic decision-making has identified several biases that violate axioms of rationality (Tsetsos, Usher, & Chater, 2010; Tversky & Kahneman, 1981). For example, adding a third option that no one would buy in the first place can still change preferences for the original options, which would go against predictions from rational choice theory in economics. Computational models that include some form of non-linear adaptation are able to reproduce the same biases (Bhatia, 2013; Roe, Busemeyer, & Townsend, 2001; Tsetsos et al., 2016). But importantly, when one assumes that the integration of information is corrupted by 'late noise' or noise in the accumulation of evidence (as opposed to noise in the initial encoding of information) (Drugowitsch, Wyart,

Devauchelle, & Koechlin, 2016), these adaptive models are more robust and in turn perform better (Spitzer et al., 2017; Tsetsos et al., 2016). Spitzer et al. (2017) showed in the same work described in the previous section through model simulations that the anti-compression of the number line in fact leads to higher accuracy when there is noise corrupting the accumulation process. These findings suggest that how structure is represented could change adaptively to the requirements of the task, but this has so far been unexplored in the structure literature.

Chapter 4 extends the results of these sequential sampling paradigms by using several numerical ranges to assess the contribution of normalization to the robustness of processing and tests the predictions from model simulations in human data. Finally, Chapter 5 uses EEG to determine whether the mechanisms described in one adaptive model, called Selective Integration (Tsetsos, Chater, & Usher, 2012), have a plausible neural foundation.

1.5 Concluding remarks

I have laid out many different ideas and models on structure representation and structure learning. This thesis does not choose a side in the debate, but rather covers topics that are relevant to all. Chapter 2 deals with the question of how the brain can abstract and generalize structure to new unseen items (p. 33). It also provides a good example of how structure representation can emerge from a simple neural network and leads to faster learning of new data with a similar structure (p. 59). Chapter 3 (p. 71) investigates how structure mapping is preserved and adapted in the brain over different contexts. Chapter 4 (p. 99) builds on the findings from Chapter 2 and 3 to ask what the most adaptive form of structure representation is in a sequential sampling task with multiple ranges. The final empirical chapter, Chapter 5

(p. 117), looks for the neural basis of a computational model that uses some form of efficient coding. Finally, in the General Discussion, Chapter 6 (p. 153), I link back to the models introduced here and discuss our findings in the light of a new theory of posterior parietal cortex.

References

- Alfred, K. L., Connolly, A. C., & Kraemer, D. J. (2018). Putting the pieces together: Generating a novel representational space through deductive reasoning. *NeuroImage*, *183*, 99–111.
- Badre, D., & Frank, M. J. (2012). Mechanisms of hierarchical reinforcement learning in cortico-striatal circuits 2: Evidence from fMRI. *Cerebral Cortex*, *22*(3), 527–536.
- Badre, D., Kayser, A. S., & D’Esposito, M. (2010). Frontal cortex and the discovery of abstract action rules. *Neuron*, *66*(2), 315–326.
- Balaguer, J., Spiers, H., Hassabis, D., & Summerfield, C. (2016). Neural mechanisms of hierarchical planning in virtual subway network. *Neuron*, *90*(4), 893–903.
- Barlow, H. B. (1961). Possible principles underlying the transformations of sensory messages. In W. Rosenblith (Ed.), *Sensory communication* (pp. 217–234).
- Behrens, T. E., Muller, T. H., Whittington, J. C., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, *100*(2), 490–509.
- Bellmund, J. L., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science (New York, N.Y.)* *362*(6415).
- Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review*, *120*(3), 522–543.
- Botvinick, M. M. (2012). Hierarchical reinforcement learning and decision making. *Current Opinion in Neurobiology*, *22*(6), 956–962.
- Bowers, J., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, *138*(3), 389–414.

- Cichy, R. M., & Pantazis, D. (2017). Multivariate pattern analysis of MEG and EEG: A comparison of representational structure in time and space. *NeuroImage*, 158, 441–454.
- Collins, A. (2018). Learning structures through reinforcement. In R. Morris, A. Bornstein, & A. Shenhav (Eds.), *Goal-directed decision making* (pp. 105–123).
- Collins, A., Cavanagh, J. F., & Frank, M. J. (2014). Human EEG uncovers latent generalizable rule structure during learning. *Journal of Neuroscience*, 34(13), 4677–4685.
- Collins, A., & Frank, M. J. (2013). Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychological Review*, 120(1), 190–229.
- Constantinescu, A., O'Reilly, J. X., & Behrens, T. E. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292), 1464–1468.
- Dabney, W., Kurth-Nelson, Z., Uchida, N., Starkweather, C. K., Hassabis, D., Munos, R., & Botvinick, M. (2020). A distributional code for value in dopamine-based reinforcement learning. *Nature*, 577(7792), 671–675.
- Dehaene, S. (2003). The neural basis of the Weber-Fechner law: A logarithmic mental number line. *Trends in Cognitive Sciences*, 7(4), 145–147.
- Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, 463(7281), 657–661.
- Doumas, L. A., Hummel, J. E., & Sandhofer, C. M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review*, 115(1), 1–43.

- Drugowitsch, J., Wyart, V., Devauchelle, A. D., & Koechlin, E. (2016). Computational precision of mental inference as critical source of human choice suboptimality. *Neuron*, *92*(6), 1398–1411.
- Flesch, T., Balaguer, J., Dekker, R., Nili, H., & Summerfield, C. (2018). Comparing continual task learning in minds and machines. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(44), E10313–E10322.
- Frank, M. J., & Badre, D. (2012). Mechanisms of hierarchical reinforcement learning in corticostriatal circuits 1: Computational analysis. *Cerebral Cortex*, *22*(3), 509–526.
- Gärdenfors, P. (2004). Conceptual spaces as a framework for knowledge representation. *Mind and Matter*, *2*(2), 9–27.
- Garvert, M. M., Dolan, R. J., & Behrens, T. E. (2017). A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *eLife*, *6*, 1–20.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.
- Gentner, D. (2003). *Why we're so smart* (D. Gentner & S. Goldin-Meadow, Eds.). Cambridge, MA: MIT Press.
- Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, *34*(5), 752–775.
- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences*, *14*(8), 357–364.
- Güçlü, U., & van Gerven, M. A. (2015). Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *Journal of Neuroscience*, *35*(27), 10005–10014.

- Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806.
- Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron*, *95*(2), 245–258.
- Holroyd, C. B., Ribas-Fernandes, J. J., Shahnazian, D., Silvetti, M., & Verguts, T. (2018). Human midcingulate cortex encodes distributed representations of task progress. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(25), 6398–6403.
- Horner, A. J., Bisby, J. A., Zotow, E., Bush, D., & Burgess, N. (2016). Grid-like processing of imagined navigation. *Current Biology*, *26*(6), 842–847.
- Jacobs, J., Weidemann, C. T., Miller, J. F., Solway, A., Burke, J. F., Wei, X. X., ... Kahana, M. J. (2013). Direct recordings of grid-like neuronal activity in human spatial navigation. *Nature Neuroscience*, *16*(9), 1188–1190.
- James, W. (1890). *The principles of psychology*. Vols. I, II.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.
- Kemp, C., & Tenenbaum, J. B. (2008). The discovery of structural form. *Proceedings of the National Academy of Sciences*, *105*(31), 10687–10692.
- Kemp, C., & Tenenbaum, J. B. (2009). Structured statistical models of inductive reasoning. *Psychological Review*, *116*(1), 20–58.
- Kemp, C., Tenenbaum, J. B., Niyogi, S., & Griffiths, T. L. (2010). A probabilistic model of theory formation. *Cognition*, *114*(2), 165–196.

- Khaligh-Razavi, S. M., & Kriegeskorte, N. (2014). Deep supervised, but not unsupervised, models may explain IT cortical representation. *PLoS Computational Biology*, *10*(11).
- Killian, N. J., Jutras, M. J., & Buffalo, E. A. (2012). A map of visual space in the primate entorhinal cortex. *Nature*, *491*(7426), 761–764.
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719.
- Kriegeskorte, N. (2015). Deep Neural Networks: A new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science*, *1*(1), 417–446.
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: Integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, *17*(8), 401–412.
- Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*(November), 1–28.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, *40*.
- Li, V., Michael, E., Balaguer, J., Herce Castañón, S., & Summerfield, C. (2018). Gain control explains the effect of distraction in human perceptual, cognitive, and economic decision making. *Proceedings of the National Academy of Sciences*, *115*(38), E8825–E8834.
- Liu, J., Li, J., Feng, L., Li, L., Tian, J., & Lee, K. (2014). Seeing Jesus in toast: Neural and behavioral correlates of face pareidolia. *Cortex*, *53*(1), 60–77.

- Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences of the United States of America*, *110*, 6139–44.
- Ma, W. J., Beck, J. M., & Pouget, A. (2008). Spiking networks for Bayesian inference and choice. *Current Opinion in Neurobiology*, *18*, 217–222.
- McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: Connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences*, *14*(8), 348–356.
- McClelland, J. L., & Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience*, *4*(4), 310–322.
- McClelland, J. L., Sadeghi, Z., & Saxe, A. M. (2016). A critique of pure hierarchy: Uncovering cross-cutting structure in a natural dataset. *Neurocomputational Models of Cognitive Development and Processing*, 51–68.
- McKenzie, S., Frank, A. J., Kinsky, N. R., Porter, B., Rivière, P. D., & Eichenbaum, H. (2014). Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron*, *83*(1), 202–215.
- Moran, R., & Tsetsos, K. (2018). The standard Bayesian model is normatively invalid for biological brains. *Behavioral and Brain Sciences*, *41*.
- Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., & Kriegeskorte, N. (2014). A toolbox for representational similarity analysis. *PLoS Computational Biology*, *10*(4).

- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1), 171–175.
- O'Keefe, J., & Nadel, L. (1980). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- O'Reilly, J. X., Jbabdi, S., & Behrens, T. E. (2012). How can a Bayesian approach inform neuroscience? *European Journal of Neuroscience*, 35, 1169–1179.
- Piaget, J. (1952). *Origin of intelligence in the child*.
- Polania, R., Woodford, M., & Ruff, C. (2019). Efficient coding of subjective value. *Nature Neuroscience*, 22(January), 134–142.
- Prelec, D. (1998). The probability weighting function. *Econometrica*, 66(3), 497–527.
- Rahnev, D. (2019). The Bayesian brain: What is it and do humans have it? *Behavioral and Brain Sciences*, 42, e238.
- Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., ... Kording, K. P. (2019). A deep learning framework for neuroscience. *Nature Neuroscience*, 22(11), 1761–1770.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108(2), 370–392.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*.
- Rumelhart, D. E., & Todd, P. M. (1993). Learning and connectionist representations. In *Attention and performance xiv: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 3–30).

- Saxe, A. M., McClelland, J. L., & Ganguli, S. (2013). Learning hierarchical category structure in deep neural networks. In M. Knauff, M. Paulen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th annual meeting of the cognitive science society* (pp. 1271–1276). Austin, TX: Cognitive Science Society.
- Saxe, A. M., McClelland, J. L., & Ganguli, S. (2019). A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 166(23), 11537–11546.
- Shafto, P., Kemp, C., Mansinghka, V., & Tenenbaum, J. B. (2011). A probabilistic model of cross-categorization. *Cognition*, 120(1), 1–25.
- Solomon, E. A., Lega, B. C., Sperling, M. R., & Kahana, M. J. (2019). Hippocampal theta codes for distances in semantic and temporal spaces. *Proceedings of the National Academy of Sciences*, 116(48), 611681.
- Spitzer, B., Waschke, L., & Summerfield, C. (2017). Selective overweighting of larger magnitudes during noisy numerical comparison. *Nature Human Behaviour*, 1, 1–8.
- Stevens, A., & Coupe, P. (1978). Distortions in judged spatial relations. *Cognitive Psychology*, 10, 422–437.
- Summerfield, C., & Li, V. (2019). Perceptual suboptimality: Bug or feature? *Behavioral and Brain Sciences*, 41.
- Summerfield, C., Luyckx, F., & Sheahan, H. (2020). Structure learning and the posterior parietal cortex. *Progress in Neurobiology*, 184, 101717.
- Summerfield, C., & Tsetsos, K. (2012). Building bridges between perceptual and economic decision-making: Neural and computational mechanisms. *Frontiers in Neuroscience*, 6(70), 1–20.

- Summerfield, C., & Tsetsos, K. (2015). Do humans make good decisions? *Trends in Cognitive Sciences*, 19(1), 27–34.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). Cambridge, MA: MIT press.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279–1285.
- Tervo, D. G. R., Tenenbaum, J. B., & Gershman, S. J. (2016). Toward the neural implementation of structure learning. *Current Opinion in Neurobiology*, 37, 99–105.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55(4), 189–208.
- Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, 109, 9659–9664.
- Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences*, 113(11), 3102–3107.
- Tsetsos, K., Usher, M., & Chater, N. (2010). Preference reversal in multiattribute choice. *Psychological Review*, 117(4), 1275–1293.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453–458.
- van Bergen, R. S., Ma, W. J., & Pratte, M. S. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience*, (October), 1–6.

- Viganò, S., & Piazza, M. (2020). Distance and direction codes underlie navigation of a novel semantic space in the human brain. *The Journal of Neuroscience*, (January), 1849–19.
- Vilares, I., & Kording, K. (2011). Bayesian models: The structure of the world, uncertainty, behavior, and the brain. *Annals of the New York Academy of Sciences*, 1224, 22–39.
- Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization. *Psychological Bulletin*, 138(6), 1172–1217.
- Walker, E. Y., Cotton, R. J., Ma, W. J., & Tolias, A. S. (2019). A neural basis of probabilistic computation in visual cortex. *bioRxiv*, 365973.
- Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron*, 81(2), 267–279.
- Yamins, D. L., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., & DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 111(23), 8619–8624.

Chapter 2

Neural structure mapping in human probabilistic reward learning

2.1 Introduction

The ability to learn rapidly from limited data is a key ingredient of human intelligence. For example, on moving to a new city, you will rapidly discover which restaurants offer good food and which neighbours provide enjoyable company. Current models of learning propose that appetitive actions towards novel stimuli are learned *tabula rasa* via reinforcement (Sutton & Barto, 2018), and these models explain the amplitude of neural signals in diverse brain regions during reward-guided choices in humans and other animals (Dolan & Dayan, 2013; O’Doherty, Dayan, Friston, Critchley, & Dolan, 2003; Schultz, Dayan, & Montague, 1997). However, model-free reinforcement learning models learn only gradually, and even when coupled with powerful function approximation methods, exhibit limited generalization beyond their

training domain (Mnih et al., 2015). This has led to the suggestion they are ill-equipped to fully describe human learning (Lake, Ullman, Tenenbaum, & Gershman, 2017).

By contrast, cognitive scientists have ascribed human intelligence to the formation of abstract knowledge representations (or “concepts”) that delimit the structural forms that new data is likely to take (Gentner, 2010; Kemp & Tenenbaum, 2008; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). Indeed, real-world data can often be described by simple relational structures, such as a tree, a grid or a ring (Tenenbaum et al., 2011). Humans may infer relational structure through probabilistic computation (Kemp, Tenenbaum, Niyogi, & Griffiths, 2010) and a long-standing theory proposes that humans understand new domains by their alignment with existing relational structures (Gentner, 1983). However, these models are often criticized for failing to specify how concepts might be plausibly encoded or computed in neural circuits (McClelland et al., 2010). A pressing concern, thus, is to provide a mechanistic account of how relational knowledge is encoded and generalized in the human brain (Tervo, Tenenbaum, & Gershman, 2016).

The current project was inspired by recent observations that the representational geometry of human neural signals evoked by symbolic numbers respects their relative cardinality. In scalp M/EEG signals, neural patterns evoked by Arabic digits vary continuously with numerical distance, such that multivariate signals for “3” are more similar to those for “4” than “5” (Spitzer, Waschke, & Summerfield, 2017; Teichmann, Grootswagers, Carlson, & Rich, 2018). Number is a symbolic system that expresses magnitude in abstract form (Bueti & Walsh, 2009; Fischer & Shaki, 2018; Walsh, 2003), and so we reasoned that continuously-varying neural signals evoked by numbers might be indexing a conceptual basis set that supports one-dimensional encoding of

novel stimuli. In the domain of reward-guided learning, a compact description of the stimulus space projects data into a single dimension that runs from “bad” to “good”. Here, thus, we asked humans to learn the reward probabilities associated with novel, high-dimensional visual images, and measured whether the stimuli come to elicit neural patterns that map onto one-dimensional neural codes for numerical magnitude.

2.2 Methods

2.2.1 Participants

Forty-nine healthy participants (Experiment 1a = 24, Experiment 1b = 25) participated in behavioural and EEG testing at the University of Oxford. Sample size for Experiment 1a was determined based on common sample sizes in the field and a similar size was used in Experiment 1b for the replication. Two participants from Experiment 1a were excluded from all analyses due to failure to learn in the bandit task (chance level performance) and one participant from Experiment 1b due to excessive movement artefacts in the recorded EEG data. All analyses were performed on the remaining 46 participants (N female = 24, N right-handed = 43, age = $24.7 \pm \text{SD } 4.5$). All participants had normal or corrected-to-normal vision, with no history of neurological or psychiatric illness. Participants were compensated for their time at a rate of £10/h plus additional bonuses based on their performance (max. £2.50 in the numerical task and £5 in the bandit task). Informed consent was given before the start of the experiment. The study was approved by the Medical Science Inter-Divisional Research Ethics Committee (R49578/RE001).

2.2.2 Experimental procedure

Both tasks (numerical and bandit) were run within a single recording session, with the order of tasks counterbalanced between participants. Stimuli were created and presented using the Psychophysics Toolbox-3 (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007) for Matlab (MathWorks) and additional custom scripts. The tasks were presented on a 20-inch screen with a resolution of 1600 x 900, at a refresh rate of 60 Hz and on a grey background. Viewing distance was fixed at approximately 62 cm. The F and J key on a standard QWERTY keyboard served as response keys for left- and right-hand responses respectively.

Numerical task

A trial started with a central dark grey fixation dot lasting 500 ms, followed by 10 Arabic digits at a rate of 3Hz (each digit was shown for 283 ms per stimulus with an inter-stimulus interval (ISI) of 67 ms; Figure 2.1A top). Numbers were drawn uniformly random from 1 to 6, with half of the stimuli randomly coloured in orange and the other half in blue. Sequence generation was unconstrained, except that the blue/orange means could not be identical. After sequence offset, participants could respond by choosing one of two response boxes on screen containing the options ‘O’ (orange) and ‘B’ (blue). Location of response options (left or right box) was fixed within the task but counterbalanced between participants. Left- and right-hand responses were used to select the left and right boxes respectively. When a response was given, the corresponding box would change to the chosen colour 100 ms before feedback. In the case of a correct answer, both letters within the boxes were replaced by a dollar sign (“\$”) accompanied by a high-pitch tone for 350 ms. Conversely, if the response was incorrect or no response was given




Figure 2.1. Task design and RSA results. **(A)** Humans performed two tasks during a single EEG recording session. In the numerical decision task, participants viewed a stream of ten digits between 1 and 6, deciding whether the blue or orange numbers had the highest/lowest average. In the bandit task, participants learned about the reward probabilities of six images (bandits) and were asked to choose between two successive bandits to obtain a fixed reward. Numbers below each frame show ISI in ms. **(B)** RSA revealed a numerical and value distance effect from 100 ms after stimulus onset (bottom coloured lines, $P_{cluster} < 0.005$). Inset shows magnitude model RDM. Shaded area represents SEM. Repeating the same analyses with correlation distance (Figure A.1, p. 172) or splitting the data per task framing group (Figure A.2, p. 173) provided highly similar results. In Figure A.3 (p. 174) we ensure that the patterns do not differ between the bandits presented first or second. **(C)** Averaged neural RDMs between 200 – 700 ms for the numerical (top) and bandit (bottom) task show a clear correspondence with the magnitude model.

within 2 seconds, the boxes would show a dash (“/”) and a low-pitch tone was played. The chosen response box remained coloured during feedback. The next trial started after an inter-trial interval (ITI) sampled uniformly between 500 and 1500 ms.

Task framing was inverted after testing the first group of participants. In Experiment 1a, participants chose the colour associated with the highest average, while in Experiment 1b they chose the colour with the lowest average. After 10 (Experiment 1a) or 20 (Experiment 1b) practice trials (excluded from analysis) each participant performed 300 trials in 6 blocks of 50 trials. Participants could take a self-timed break in between blocks.

Bandit task

Bandits were represented by six unique coloured drawings of donkeys, freely available on the internet, and each donkey was coloured differently (green, purple, orange, yellow, red and blue) using GIMP (<https://www.gimp.org>) and superimposed on a light grey background (Figure 2.1 A bottom). Each image was associated with one of six stationary reward probabilities, linearly spaced between 0.05 and 0.95 [0.05, 0.23, 0.41, 0.59, 0.77, 0.95], and these were assigned randomly to donkeys for each participant.

The experiment consisted of two learning phases and one test phase. In the first learning phase (L1), all six bandits were presented on screen simultaneously, in a 2 x 3 configuration. Participants could click on each bandit 36 times, in any order they preferred. Every time a bandit was selected, the image was replaced by a feedback sign and the surrounding frame adopted the colour of the chosen bandit. Feedback in a successful trial consisted of a centrally presented dollar sign and a high-pitch tone, otherwise a dash appeared paired with a low-pitch tone. Feedback was not drawn

probabilistically in this phase: the number of successful trials was determined based on the maximum number of clicks per bandit and its associated reward probability. After a bandit was chosen 36 times, it was masked and remained unavailable until the end of the phase.

In the second learning phase (L2), participants were presented with two random bandits sequentially, identical to the test phase (see below). Phase L2 only differed from the test phase in that it included 50% pseudo-choice trials, where the computer determined which bandit had to be chosen, in order to encourage exploration of all bandits. During these trials, the computer-chosen bandit and corresponding response key were highlighted with a brown frame. All pseudo-choice bandits were assigned equally often and balanced over presentation order. It was emphasized to participants that the pseudo-choice bandits were selected randomly. Phase L2 lasted for 2 blocks of 60 trials with a short break in between blocks.

Each trial in phase L2 and in the test phase started with the presentation of a white fixation dot for 500 ms. The fixation dot would disappear 250 ms before presentation of the first bandit. Two bandits were then presented sequentially, each for 500ms with an ISI of 250 ms. Afterwards, participants had to choose which bandit they preferred. Two response boxes on screen indicated the choice options: ‘A’ referring to the first bandit presented and ‘B’ to the second. The location of ‘A’ and ‘B’ (left or right box) alternated randomly from trial to trial. The chosen box would then change to the colour of the chosen bandit 200 ms before feedback was given. Reward was determined randomly according to the reward probability of the chosen bandit. If the choice was successful, both response boxes would contain a dollar sign and a high-pitch tone played for 500 ms. If the choice entailed no reward or no response was given within 2 seconds, both boxes would show a

dash and a low-pitch tone was played. The fixation dot also turned red when no response was given. A new trial started after an ITI of 500 ms. The test phase consisted of 10 blocks of 60 trials. By the probabilistic nature of the bandit outcomes, the subjective ranking of the bandits could potentially vary over the course of the experiment. It was therefore emphasized to participants before the test phase that the actual reward probabilities of the bandits would never change.

2.2.3 EEG acquisition

The EEG signal was recorded using 61 Ag/AgCl sintered surface electrodes (EasyCap, Herrsching, German), a NeuroScan SynAmps RT amplifier, and Curry 7 software (Compumedics NeuroScan, Charlotte, NC). Electrodes were placed according to the extended international 10-20 system, with the right mastoid as recording reference and channel AFz as ground. Additional bipolar electrooculography (EOG) was recorded, with two electrodes placed on either temple for recording horizontal EOG and two electrodes above and below the right eye for vertical EOG. All data was recorded at 1 kHz and low-pass filtered online at 200 Hz. All impedances were kept below 10-15 k Ω during the experiment.

2.2.4 EEG pre-processing

The data from both tasks were pre-processed following the same pipeline, using functions from the EEGLAB toolbox (Delorme & Makeig, 2004) for Matlab and custom scripts. First the data were down-sampled to 250 Hz, low-pass filtered at 40 Hz and then high-pass filtered at 1 Hz. The continuous recording was visually screened for excessively noisy channels and these were interpolated by the weighted average of the surrounding electrodes. The

data was then offline re-referenced to average reference. In the numerical task, epochs were extracted from 1 second before fixation dot onset to 5.5 seconds after. In the bandit task, epochs were extracted from 0.5 seconds before fixation dot onset to 3 seconds after. Epochs were baselined relative to the full pre-fixation time window. Epochs containing atypical noise (such as muscle activity) were rejected after visual inspection. We then performed Independent Component Analysis (ICA) and removed components related to eye blink activity and other artefacts (manually selected for each participant). Lastly, the trial epochs were split into smaller stimulus epochs for each digit/bandit and re-baselined based on the pre-stimulus onset time window. For the numerical task, these epochs spanned -65 ms to 850 ms relative to stimulus onset. For the bandit task, they spanned -250 ms to 750 ms. Final analyses focused on the overlapping time window of -65 to 750 ms in both tasks.

2.2.5 Statistical procedure

Experiment 1b was pre-registered as a control experiment (DOI 10.17605/osf.io/ym3gu), directly replicating the bandit task and inverting the task framing for the numerical task. We confirmed our proposed hypothesis that task framing would not affect either the direction of the univariate parietal effects (Figure A.6, p. 177) or the multivariate patterns (Figure A.2, p. 173). To increase the power of our analyses, we collapsed the two data sets. All analyses were conducted on the subject level and statistics are reported for the group level.

2.2.6 Delta-rule model

We estimated subjective probabilities for each bandit using a delta-rule model. On every trial, the model compares the (subjective) values of the two offered bandits and updates the value of the chosen bandit based on the observed reward during the task:

$$V_c(t+1) = V_c(t) + \alpha[R(t) - V_c(t)] \quad (2.1)$$

where $V_c(t)$ is the value of the bandit chosen by the participant on trial t , α the learning rate and R the received reward (either 0 or 1). Value of the chosen bandit is updated for the next trial $t+1$, by taking the difference between the observed reward $R(t)$ and the expected reward $V_c(t)$ modulated by the size of the learning rate. To obtain estimates of the model's choices, the values of the two bandits are passed through a sigmoidal response function:

$$p(A) = \frac{\lambda + (1 - 2\lambda)}{1 + e^{-\Delta V/s}} \quad (2.2)$$

where $\Delta V = V_A - V_B$, i.e. the difference in value between the first (A) and second (B) bandit. The policy parameters s and λ indicate the slope and termination point (lapse rate) of the logistic choice function. The latter was fixed to 0.05, equivalent to the bounds of the reward probabilities in the task, and the parameters α and s (learning rate and slope) were fit to the data. Best-fitting parameter estimates were obtained by minimizing the negative log-likelihood function using optimization tools in Matlab. Bandit probabilities in the test phase were initialized according to the estimates of subjective reward probability obtained from free-choice trials from phase L2. Search space was restricted for both parameters between 0.0001 and 0.5. Best fitting parameters were then used to estimate trial-by-trial subjective ranks

for all bandits for each participant. per participant by classifying the bandits according to their subjective values.

2.2.7 Representational Similarity Analysis (RSA)

The pre-processed EEG data was first z-scored over all trials, per electrode and time point. To obtain the condition-specific activations at each time point and electrode, we constructed a design matrix using dummy coding for each condition within a task (numbers 1 to 6 and the six bandits) and subsequently estimated beta coefficients using a linear regression model for each number or (ranked) bandit. These beta coefficients then reflected the trial-average response per condition at each time point and electrode. We then calculated the Euclidean (and correlation) distance between the whole-scalp neural signals of each condition pair (e.g. bandit 1 and bandit 2), resulting in a 6 x 6 representational dissimilarity matrix (neural RDM) at each time point for both tasks separately (Figure 2.1 C). To test for patterns of numerical distance in the neural RDMs, we created a 6 x 6 magnitude model RDM in which the predicted dissimilarity linearly increased from 0 to 1 as a function of the numerical difference between two numbers. The upper triangles of the model and neural RDM were subsequently correlated using Kendall τ -a rank correlation (Figure 2.1 B) (Kriegeskorte & Kievit, 2013; Nili et al., 2014).

In the cross-validation analyses, we followed a similar pipeline as the within-task RSA described above, except that beta estimates of condition activation from both tasks were concatenated before calculating the Euclidean distance between all 12 conditions, resulting in a 12 x 12 RDM containing both within-task (e.g. number-number) and between-task (e.g. number-bandit) dissimilarities (Figure 2.2 A). To account for potential differences in the time at which a magnitude code was decodable, we conducted this analysis for all

possible combinations of time points (cross-temporal RSA) (King & Dehaene, 2014). Neural RDMs were smoothed over time through convolution with a 60 (x 60) ms uniform kernel. In cross-validation, model RDMs were correlated with the lower rectangle of the neural RDM, containing the between-task dissimilarities. The diagonal of the rectangle was included, since on-diagonal information is non-redundant in cross-validation. Correlations were baseline corrected by subtracting the average correlation in the pre-stimulus period to not bias our cluster-identifying algorithm. Significant clusters were identified using cluster-corrected non-parametric permutation tests (iterations = 1000, cluster-defining and cluster-level thresholds at $P < 0.005$, unless stated otherwise) (Maris & Oostenveld, 2007).

We also tested to what extent the choice behaviour of a task was reflected in the neural patterns through multiple regression (Figure 2.4 A-D). For the numerical task, a model RDM was calculated for each participant based on the differences in choice probability for each number pair. The bandit model RDM was constructed by taking the average probability of choosing the most valuable bandit (according to the delta-rule model estimates) for any bandit pair. Both models were vectorized and z-scored and entered as two regressors in a model explaining variance in the neural RDM of either the numerical task (Figure 2.4 C) or the bandit task (Figure 2.2 D). Similarly, for the CPP control analyses (Figure A.7 D-E, p. 178), a model RDM was constructed based on the differences in average peak activity per condition (see below). The CPP model was then entered in a competitive regression together with the magnitude model, to determine whether the variance in neural patterns could be explained by a magnitude code over and above the univariate findings.

2.2.8 CPP analysis

Normalized EEG epochs were averaged for each digit (independent of colour category) or for each subjective bandit rank. Based on previous research (Spitzer et al., 2017; Twomey, Murphy, Kelly, & Connell, 2015), we selected seven centro-parietal electrodes (CP1, P1, POz, CPz, CP2 and P2) and averaged the event-related potentials (ERP) over these electrodes for each stimulus. Next, we sought to identify for each task the time window where the disparity in ERP signals was greatest between stimulus types, using a non-parametric omnibus test (Kruskal-Wallis test) at every time point. To avoid circular inference, the test was performed in a leave-one-out fashion, determining significance based on the remaining 45 participants. The largest cluster of adjacent significant time points ($P < 0.01$) was then determined using FDR correction for multiple comparisons (Benjamini & Hochberg, 2009). Activity within each individual’s time window was then averaged for each condition and differences in the ERP averages were taken to construct RDMs that were used as control models for RSA (Figure A.7, p. 178).

2.2.9 Dimensionality reduction

We assessed the dimensionality of the neural data through Singular Value Decomposition (SVD), a method that allows to efficiently obtain principal components in the data through linear transformation. After estimating beta coefficients for each condition, we used SVD at each time point to obtain the diagonal matrix Σ that contained the six singular values. We then systematically reduced the dimensionality of the data by removing the last column of Σ , reconstructed the data under the reduced dimensionality and followed the rest of the cross-validation pipeline as described above. Changes in the strength of the cross-validation were tested by comparing the model –

EEG RDM correlations averaged over the 350-600 ms time window, under the different dimensionalities of the data (Figure 2.5 C).

2.2.10 Psychometric model

To estimate numerical magnitude distortion in the numerical task, we fitted a psychometric model to the choice data (Figure 2.3 F). The model is adapted from that described in Spitzer et al. (2017). First, input values (digits numbers 1 – 6 in the task) are normalized between 0 and 1 in six equidistant steps. These normalized values nX are then transformed into a subjective decisional value $\hat{X}$ by exponentiating them with free parameter k .

$$\hat{X} = nX^k \quad (2.3)$$

When $k = 1$, $\hat{X}$ is equal to nX . When $k > 1$, decision values are exponential, giving relatively higher weights to larger numbers and being more indifferent to smaller numbers. Conversely, when $k < 1$, decision values are compressed, giving lower weight to the larger numbers. Next, we model the trial-level decision-value as the sum over all sample $\hat{X}$.

$$DV = \sum_{i=1}^{10} \hat{X}_i \times c_i \times l^{10-i} \quad (2.4)$$

where c_i is an indicator variable that codes for stimulus category (orange = -1, blue = 1) and l is a leakage term that exponentially discounts earlier samples in the stream. Finally, the model uses a logistic function for computing choice probabilities:

$$p(blue) = f - \frac{1}{1 + e^{-(B+DV)/s'}} \quad (2.5)$$

where $p(blue)$ is the probability of choosing the blue category. s' is the

inverse slope of the logistic choice function and B captures a simple response bias towards one of the two categories. Finally, f is an indicator variable that codes for the framing of the task, inverting the choice probability for the low frame ($f = 1$) compared to the high frame ($f = 0$). Best fitting parameters for k [0.01 – 10], l [0 – 1], B [unconstrained] and s' [0.001 – 8] were obtained using standard optimization tools in Matlab.

Model-free decision weights were estimated for each participant using an averaging approach to compare to the model predictions. Choice probabilities for each digit were calculated by averaging over all responses to trials where that digit was present.

2.2.11 Neurometric mapping

In order to obtain an estimate of the distortion in neural representational geometry, we generated a range of candidate model RDMs computed from features that were distorted by a parameter k , analogous to the psychometric model. Six equidistant values between 0 and 1 were raised to the power of k (Equation 2.3) and model RDMs were constructed based on the Euclidean distances after distortion. Parameters that fitted the neural data best were found through exhaustive search through values of k from 0.35 – 3 in steps of 0.01, iteratively computing the Kendall τ -a correlation between the upper triangle of the model RDM and the neural RDM averaged over a time window between 350 – 600 ms. The distortion k associated with the highest correlation was determined in each participant individually and subjected to group-level analysis. Representational geometries for the two tasks were stable in this time window based on the analysis depicted in Figure 2.2 B.

2.2.12 Neural network

We constructed a simple feedforward neural network with 21 input units (1 bias unit and two input modules of 20 units, X_A and X_B), 10 hidden units H and 10 output units Y (Figure 2.6 A). The network was trained (learning rate = 0.001) to map inputs x_a and x_b , both consisting of 6 stimuli with 20 features each, onto 6 random vectors. Both inputs were generated by drawing random values from a standard normal distribution. In the crucial test condition, both x_a and x_b shared a second-order structure of gradually increasing dissimilarity. This was achieved by successively flipping the sign of two more features for each adjacent stimulus. The network learned to minimize the cost-function (cross-entropy) with respect to the supervision signal and via backpropagation, first on inputs x_a arriving at input units X_A , and after 106 iterations it was then retrained on inputs x_b fed into units X_B .

Two control conditions were included to assess the contribution of shared structure to retraining performance. In one control, the model was initially trained on an input x_a that was entirely random, i.e. not constructed using the sign-flipping method. In the second control, hidden-to-output weights W_2 were shuffled before retraining on x_b . Simulations were run 100 times for each of the conditions (x_a and x_b shared structure vs. x_a was random; W_2 unshuffled vs. shuffled).

Next, RSA was performed on the hidden unit activations after convergence had been achieved on retraining to test for shared representational structure in the hidden units. Each stimulus from x_a was fed into X_A and x_b into X_B to obtain hidden unit activations for each stimulus. A cross-validation RDM was then constructed based on the differences in activations between x_a and x_b for each stimulus and correlated with a magnitude model RDM (Figure 2.6 E).

2.3 Results

Whilst undergoing scalp EEG recordings, human participants ($N = 46$) completed two tasks: a numerical decision task and a probabilistic reward-guided learning task (Figure 2.1 A). In the numerical task participants viewed rapid streams of ten Arabic digits (1 to 6) and reported whether numbers in orange or blue font had the higher (Experiment 1a, $N = 22$) or lower (Experiment 1b, $N = 24$) average. The reward-learning task was based on the multi-armed bandit paradigm that has been used ubiquitously to study value-guided decision-making (Dolan & Dayan, 2013). Participants learned the reward probabilities associated with six unique novel images (coloured donkeys), which paid out a fixed reward with a stationary probability (range 0.05-0.95). These probability values were never signalled to the participant but instead acquired by trial and error in an initial learning phase. In the test phase, we asked participants to decide between two successive donkeys to obtain a reward, and estimated trial-wise subjective probability estimates for each bandit by fitting a delta rule model to choices (Sutton & Barto, 2018). Throughout these phases, the bandits were never associated with numbers in any way.

2.3.1 Shared magnitude representation for numbers and probabilistic rewards

Using representational similarity analysis (RSA) (Kriegeskorte & Kievit, 2013) we replicated the previous finding (Spitzer et al., 2017; Teichmann et al., 2018) that patterns of neural activity across the scalp from ~ 100 ms onwards were increasingly dissimilar for numbers with more divergent magnitude, i.e. codes for “3” and “5” were more dissimilar than those for “3” and “4” (Figure 2.1 B, green line). This occurred irrespective of task framing (report higher

vs. lower average) and category (orange vs. blue numbers), suggesting that neural signals encoded an abstract representation of magnitude and not solely a decision-related quantity such as choice certainty (Spitzer et al., 2017). Next, we used RSA to examine the neural patterns evoked by bandits. We found that multivariate EEG signals varied with subjective bandit ranks, with bandits that paid out with nearby probabilities eliciting more similar neural patterns (from 100 ms onwards; Figure 2.1 B, blue line).




Figure 2.2. Cross-temporal cross-validation RSA. **(A)** Averaged neural RDM from 350-600 ms across the numerical (1-6) and bandit (b1-b6) task. Upper left and lower right quadrants show representational dissimilarity for numbers (‘N’) and bandits (‘B’) respectively, i.e. within-task RDM; lower left/upper right quadrants show cross-validated dissimilarity between numbers and bandits (‘N-B’), i.e. between-task RDM. **(B)** Cross-temporal cross-validated RSA revealed a stable magnitude representation that was shared between the two tasks ($P_{cluster} < 0.005$) around 350 – 600 ms. **(C)** To ensure that our cross-validation was not driven by a single number/bandit pair, we systematically removed each number/bandit pair and repeated the cross-temporal cross-validated RSA on the subset data. Each data point can range from 0 (no significant cross-validation for any number/bandit pair) to 6 (cross-validation always reached significance). Between 400 – 600 ms our cross-validation was robust to the removal of any number/bandit pair (Figure A.4 (p. 175) shows the results isolated when excluding number/bandit pair 6).

Our key question was whether there was a shared neural code for numerical magnitude and reward probability. We found that EEG signals elicited by digit “6” were more similar to those evoked by the most valuable bandit, and digit 1 predicted the bandit least likely to pay out, with a similar convergence for intermediate numbers and bandits (Figure 2.2 A). Cross-validation of neural signals elicited by all numbers (1 to 6) and bandits (inverse ranks 1-6) was stable and reliable from 300 - 650 ms post-stimulus, as demonstrated by cross-temporal RSA (Figure 2.2 B). We conducted several control analyses to further explore the nature of this effect. The cross-validation effect was not driven by patterns within a single number/bandit pair, as it remained robust to the removal of any one of the six number/bandit pairs (Figure 2.2 C). In particular, although the number/bandit 6 cross-validation appeared more dissimilar to other numbers/bandits, the effect persisted when only numbers/bandits 1-5 were included in the analysis (Figure A.4, p. 175). We then asked whether each number was more similar to its equivalent bandit than to other bandits (e.g. number 3 and bandit 3, the “on-diagonal” information in Figure 2.3 A), by computing an “Exemplar Discriminability Index” (EDI) (Nili, Walther, Alink, & Kriegeskorte, 2016). Additionally, we asked whether numbers showed a gradually increasing dissimilarity to non-identical bandits (e.g. whether number 3 was more similar to bandit 2/4 than bandit 1/5, the “off-diagonal” information in Figure 2.3 B). Both of these effects were independently reliable, suggesting not only that each number shares a representation with its corresponding bandit, but that the transitive patterns of encoding numbers and bandits are in a common register in neural signals.




Figure 2.3. Detailed examination of the magnitude pattern in cross-validation. **(A)** The presence of a one-to-one mapping along the number line (i.e. number 1 to least valued bandit) was tested using the ‘Exemplar Discriminability Index’ (EDI), a measure that indicates how much better items on average map onto iterations of the same item compared to different items. The EDI is calculated by subtracting the mean on-diagonal distances from the mean off-diagonal distances (top illustration). Significant exemplar discriminability, where numbers most closely resembled their equivalent bandit, arose at the same time as our main findings. **(B)** Removing the diagonal from the lower rectangle in the cross-validation RDM (i.e. the dissimilarity between distributed responses to corresponding stimuli in the two tasks) reproduced the results from our main analysis, suggesting the effect was not mainly driven by matching stimulus pairs (e.g. number 6 and most valuable bandit), but by a gradual distance effect ($P_{cluster} < 0.005$).

2.3.2 Relating numerical magnitude representation to choice behaviour

Past work has identified overlapping choice biases in numerical cognition and economic decisions (Kanayet, Opfer, & Cunningham, 2014; Schley & Peters, 2014). Next thus, we asked how patterns of behaviour in the number and bandit tasks were related to one another by creating choice matrices encoding the difference in relative weight given to each number or bandit in the choices made by participants. For the numerical task, we computed decision weights for each number in the choice using an averaging approach and plotted the relative difference in these weights for each combination of numbers (Figure 2.4 A). For the bandit task, this was simply the probability of responding “highest” for each combination of bandits (Figure 2.4 B). Although choice matrices for numbers and bandits were on average correlated across participants ($\bar{r}_\tau = 0.32$, $Z = 5.65$, $P < 0.0001$, Wilcoxon signed rank test), this correlation disappeared when subtracting the group average choice RDMs and correlating the residual RDMs ($\bar{r}_\tau = -0.02$, $Z = -0.49$, $P = 0.63$). In other words, we were not able to identify shared variation in individual weighting of numbers 1-6 and bandits 1-6 in behaviour alone, perhaps because of the different nature of the decision required in each task.

Subsequently, we asked how these choice matrices explained variance in patterns of neural similarity, i.e. whether the behaviour explained shared variation in the neural structure alignment for numbers and bandits. To this end, we substituted the linear representational distance matrix used for earlier analyses (i.e. one that assumes equal spacing among adjacent numbers and bandits) with the subject-specific choice matrices computed from behaviour and repeated the analyses above. This allowed us to ask how patterns of neural similarity among both numbers and bandits were




Figure 2.4. Behavioural analyses. **(A)** Group-averaged difference RDM of decision weights in the numerical task. We calculated the weight of each number on participant's choice, independent of colour, and created a choice RDM through pairwise differences of all numbers. **(B)** Group-averaged choice RDM in the bandit task. Each cell contains the probability of choosing the highest valued bandit on trials where bandit x and bandit y were presented together. Participants never encountered trials where the same bandit was presented twice.

(C) Both choice matrices were inserted in a multiple regression explaining EEG patterns to establish a potential link between behavior and neural patterns in either task. Only the numerical choice RDM explained neural patterns of the numerical task. (D) In contrast, both choice RDMs significantly explained neural patterns in the bandit task, possibly indicating that participants relied on a more general understanding of numbers in both tasks (bottom colored lines, $P_{cluster} < 0.005$). Shaded area represents SEM. (E) Distributions of $\log(k)$ after fitting a power-law model of the form x^k to the decision weights in the numerical task (yellow) and the average neural patterns in the numerical (green) and bandit task (blue) between 350 and 600 ms. In all cases the best fitting parameter $\log(k)$ was significantly greater than 0 (linear), indicating an overweighting/increasing dissimilarity for larger quantities (numbers 5-6 or highest valued bandits). (F) Decision weights under the median estimated k for psychometric (yellow) and neurometric (green/blue) fits, compared to the median true human decision weights in the numerical task.

explained by variance in subject-specific choice matrices. For each neural pattern, we used a regression-based approach in which choice matrices from numbers and bandits were entered as competing regressors. Interestingly, we found that choice patterns from the numerical task explained variance in the neural patterns for both numbers and bandits, but choice patterns from the bandit task only explained variance in the neural patterns of the bandit task (Figure 2.4 C-D) ($P_{cluster} < 0.005$). One interpretation of this finding is that humans used their intrinsic sense of magnitude when forming neural representations of the bandits, but not vice versa. We note in passing that this asymmetry is not secondary to the ordering of the two tasks, which was fully counterbalanced across participants.

In previous work, we observed that participants tended to give relatively greater weight to larger magnitudes during the numerical decision

task, e.g. numbers “5” and “6” had disproportionate impact on averaging judgments (Spitzer et al., 2017). This finding was replicated in the current data (Figure 2.4 E-F). Human choices were best fit by a power-law model in which participants averaged and compared distorted numerical values x^k with $k = 2.04 \pm 1.11 (k > 1 : t(45) = 12.47, P < 0.001)$. This prompted us to ask whether any shared variance between behavioural choice matrices and neural signals for the two tasks could be explained by subject-specific differences in the pattern of compression or anti-compression in the mental number line, as characterized by this model. Turning to the neural data, we thus generated candidate representational dissimilarity matrices (RDMs) under the assumption that distance in neural space can likewise be non-linear and best described by a distortion given by the same power-law model, i.e. of the form x^k . We found that in both numerical and bandit tasks the best fitting RDM was parametrized by $k > 1$ [*numerical* : $k = 1.73 \pm 0.82, t(45) = 14.34, P < 0.001$; *bandit* : $k = 1.72 \pm 0.85, t(45) = 13.65, P < 0.001$]. In other words, we observed that the anti-compressed number line estimated from behaviour was reflected, on average, in both the neural representation of numbers and bandits. However, when we correlated the subject-specific model parameter k from behaviour with estimates obtained from the neural data from either task, we found that although the degree of behavioural anti-compression strongly predicted the neural anti-compression for the numerical task ($\rho = 0.57, P = 0.0004$), it did not for the bandit task ($\rho = -0.13, P = 0.39$). One interpretation of this result is that the variance linking the mental number line to the representation of bandits (i.e. from Figure 2.4 D) is not simply due to individual differences in compression or anti-compression of the mental representations of numbers and bandits but must lie in a subspace not captured by this simple unidimensional model.

2.3.3 Dimensionality of magnitude representation

Recent work has suggested that during categorization, posterior parietal neurons in the monkeys are strikingly low-dimensional, as if the parietal cortex were engaging in a gain control process that projected stimulus features or timings on a single axis (Fitzgerald et al., 2013; Ganguli et al., 2008; Platt & Glimcher, 1999; Wang, Narain, Hosseini, & Jazayeri, 2018). Indeed, we observed a centro-parietal positivity (CPP) that varied with the magnitude of both numbers and reward probabilities (Figure 2.5 A-B). This signal resembles a previously described EEG signal, that has been found to scale with the choice certainty in perceptual (O’Connell, Dockree, & Kelly, 2012) and economic tasks (Pisauro, Fouragnan, Retzler, & Philiastides, 2017). However, in our numerical task the CPP followed an approximately ascending pattern from lower to higher numbers regardless of task framing (Figure A.6, p. 177) or colour category. The cross-validation effect persisted even after the CPP had been regressed out of the data (Figure A.7, p. 178). This suggests that (a) the CPP in our task may represent a notion of magnitude, not a certainty signal alone; and (b) that this signal is not the sole driver of our multivariate findings.

Nevertheless, to understand the dimensionality of the number and bandit representations (and the subspace in which they aligned), we used two dimensionality reduction techniques, singular value decomposition (SVD) and multidimensional scaling (MDS). First, using SVD, we systematically removed dimensions from the EEG data and recomputed our number-bandit cross-validation scores (Figure 2.5 C). We found that probabilistic reward learning was supported by a low-dimensional neural magnitude code, with reliable effects persisting when all but 2 eigenvectors were removed from the data but significantly attenuated when only a single dimension was retained




Figure 2.5. Dimensionality of magnitude representation. **(A)** Regression coefficients associated with numbers 1-6, independent of task framing (report highest/lowest) or category (blue/orange) at highlighted centro-parietal electrodes. Grey shaded area shows time of greatest disparity between signals (Kruskal-Wallis, $P_{FDR} < 0.01$). Scalp map inset shows response amplitude for number 6 during identified time window. Coloured shading represents SEM. The ascending direction of the univariate responses was independent of task framing (Figure A.6, p. 177) **(B)** Equivalent analysis for bandits b1 (lowest value) to b6 (highest value) in the bandit task. Scalp map shows response amplitude for highest subjectively valued bandit b6. **(C)** Dimensionality of the data was iteratively reduced using SVD and the strength of cross-validation under each new dimensionality was assessed by comparing the mean cross-validation in the 350 – 600 ms time window (bottom plot). Each cell in the grid contains the t- and p-value of a pairwise comparison of mean CV under different dimensionalities of the data. Reduction to one (and to a lesser degree two) dimension(s) significantly reduced the size of the effect. **(D)** Multidimensional scaling (MDS) revealed two principal axes that describe the data: a magnitude axis approximately following the number/bandit order and a certainty axis distinguishing inlying (e.g. 3,4) from outlying (e.g. 1,6) numbers or bandits.

in the EEG data. Indeed, 3D and full (high-dimensional) solutions led to statistically equivalent cross-validation, with some attenuation of the effect when only 2 dimensions were retained and a more dramatic decrease with only a single dimension. To further establish that the cross-validation effect was not solely driven by the observed univariate activity, we again used SVD to remove the first dimension and re-computed the cross-validation statistics (Figure A.8, p. 179). A significant cluster of cross-validation emerged at the same time as the originally observed effect together with a previously unobserved cluster later in time ($P_{cluster} < 0.05$). In summary, aside from a major univariate component to our cross-validation effect, there remains a shared pattern that lies in higher dimensions of the data.

Secondly, we used MDS to visualize the first dimensions of the concatenated number/bandit data. This disclosed an axis pertaining to magnitude and another approximately corresponding to certainty along which, especially for the bandits, the large (or best) and small (or worst) items diverged from the others (Figure 2.5 D). In other words, the numbers and bandits align principally along a single magnitude axis but with an additional contribution from a second factor potentially encoding choice certainty.

2.3.4 Neural network simulations

What are the potential benefits of the shared coding scheme we observed in neural signals? One possibility is that shared structure can promote generalization, such that new relational structures (i.e. the transitive relations among bandits as a function of their reward probability) are learned faster and more effectively when an existing scaffold (such as a transitive representation of number) has been previously learned. We are unable to test for this benefit directly in our human data, because all participants were numerate adult hu-




Figure 2.6. Neural network simulations. **(A)** Schematic depiction of network structure and training. The network was first trained to classify inputs x_a fed into units X_A (lower green circles). Inputs x_a consisted of 6 stimuli that either exhibited gradual increasing dissimilarity or were shuffled as a control (stimulus RDMs shown next to examples of x_a and x_b). After convergence, the model was trained on new input x_b that were fed into a separate input stream X_B (lower blue circles). Inputs x_b were different to x_a but exhibited the same similarity structure. **(B)** Loss plotted over the course of training (left panel) and retraining (right panel) for the test (red) and shuffled control (blue) conditions. Learning was faster during training for control stimuli, but retraining was faster when x_a and x_b exhibited shared similarity structure. **(C)** Loss for control simulations where hidden-to-output weights $W2$ were shuffled between training and retraining, suggesting successful transfer depends on structure encoded in $W2$. **(D)** Mean loss for first 1000 cycles after retraining. **(E)** Cross-validation RSA on hidden unit activation for all stimuli in x_a and x_b after retraining. Hidden unit activations exhibit shared similarity structure only when $W2$ remains unshuffled and x_a and x_b share structure. Error bars show SEM over 100 network simulations with different initialization.

mans, denying us an appropriate control condition. However, to demonstrate the theoretical benefit of shared coding at the mechanistic level, we turned to a simple computational tool, a feedforward neural network (Figure 2.6 A). Neural networks are not constrained to make inferences over structure, but structured representations may emerge naturally in the weights during training (McClelland et al., 2010). To explore how these representations might be learned and generalized at the mechanistic level, we turned to a simple computational tool, a feedforward neural network. Neural networks are not constrained to make inferences over structure, but structured representations may emerge naturally in the weights during training (McClelland et al., 2010). Here we confronted the network with two different stimulus sets in turn that (like our numbers and bandits) shared the same similarity structure. We then asked if the shared structure facilitates retraining on the second set after learning the first. The network was first trained on inputs x_a arriving at input units X_A , and after convergence, retrained on inputs x_b fed into units X_B (where X_A and X_B are separate input modules that project to a common hidden layer H). Inputs x_b were 6 random vectors constructed to have the same continuously-varying similarity structure as the bandits, whereas inputs x_a consisted of either a different set of 6 random vectors with the same second-order structure, or a shuffled control lacking the second-order structure. Relearning on x_b proceeded faster when inputs shared a common structure with x_a (Figure 2.6 B-D). In a second control we shuffled the weights W_2 connecting the hidden layer to the output layer after convergence on input x_a , destroying the mapping of activity patterns in the hidden layer to the output layer. RSA conducted after retraining revealed reliable cross-validated patterns of activity in the hidden units only for the condition where x_a and x_b shared an underlying structure and weights W_2

were kept intact alone (Figure 2.6 E), mirroring the result from the human neural data.

2.4 Discussion

We report that during a probabilistic reward-guided learning task involving arbitrary images (‘bandits’), stimuli with high payout probability shared a neural code with larger numbers, and those with lower value shared a neural code with lower numbers. We interpret these data as indicating that an abstract neural code for magnitude forms a conceptual basis set or “scaffold” for learning new information, such as the reward probabilities associated with novel stimuli. Rather than encoding stimulus value in an unstructured value function or lookup table (as is common in RL models), our data suggest that humans project available stimuli onto a low-dimensional axis that runs from “bad” to “good”. This neural axis is aligned with the mental number line, suggesting that humans recycle an abstract concept of magnitude to encode reward probabilities.

Learning a structured representation of value will have the benefit of allowing new inductive inferences, such as inferring transitive preferences among economic goods without exhaustive pairwise comparison (Alfred, Connolly, & Kraemer, 2018), and facilitate read-out in downstream brain areas, related perhaps to the notion of a ‘common currency’ for reward (Levy & Glimcher, 2012).

Our neural network simulations provide a demonstration of how shared structure can promote generalization and thus faster learning, even for stimuli with different physical input features. However, our design was not suited to directly test a benefit of learning between tasks, since our participants were numerate adults when they entered the experiment. In further work, it would

be of particular interest to teach participants two new transitive structures, each associated with a different stimulus set, and test whether the extent to which the neural codes align predicts learning rates for the second stimulus set – a direct prediction that emerges from our computational simulations.

We did however find a link between the transitive neural codes and participants' behaviour. Choice patterns in the two tasks were positively correlated, and how participants treated numbers in the numerical task was reflected in the neural patterns of both tasks. Furthermore, we found that participants tended to overweight larger numbers and this non-linear weighting correlated highly with non-linear representations at the neural level, at least in the numerical task. Why we did not find a correlation with non-linearities in the neural patterns of the bandit task is unclear. One speculation is that the non-linearity arises at a later stage in the processing of stimuli and is specific to the task at hand.

One major limitation of our approach is the limited spatial resolution of EEG. This leaves open the question of the true dimensionality of the shared neural code. Our investigations using dimensionality reduction techniques indicate that the code in EEG signals is low-dimensional but not simply univariate. Thus, whilst our work is consistent with previous studies showing that the amplitude of centro-parietal EEG signals scales with number (Spitzer et al., 2017) and the value of economic prospects, such as food items (Pisauro et al., 2017), it also suggests a more complex pattern encoding the shared structure among stimuli defined by transitive relations. However, it would be instructive to measure the effect using techniques that potentially afford higher spatial resolution, such as MEG or fMRI. More generally, however, our work is consistent with theories that have suggested that neural signals for magnitude in the parietal cortex may provide a conceptual bridge between

different metrics such as space, time and number (Buetti & Walsh, 2009; Chafee, 2013; Fitzgerald et al., 2013; Parkinson, Liu, & Wheatley, 2014; Walsh, 2003).

Data and code availability

This work has been published in *eLife*: <https://doi.org/10.7554/eLife.42816.001>. All code and materials to reproduce the analyses and experiments are available at https://github.com/summerfieldlab/Luyckx_etal_2019. Data to reproduce the results are available from the Dryad Digital Repository: <https://doi.org/10.5061/dryad.7k7s800>. Raw EEG data files are available upon request.

References

- Alfred, K. L., Connolly, A. C., & Kraemer, D. J. (2018). Putting the pieces together: Generating a novel representational space through deductive reasoning. *NeuroImage*, 183, 99–111.
- Benjamini, Y., & Hochberg, Y. (2009). Controlling the False Discovery Rate : A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1), 289–300.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, (10), 433–436.
- Bueti, D., & Walsh, V. (2009). The parietal cortex and the representation of time, space, number and other magnitudes. *Philosophical Transactions of the Royal Society B*, 364(1525), 1831–1840.
- Chafee, M. V. (2013). A scalar neural code for categories in parietal cortex: Representing cognitive variables as "more" or "less". *Neuron*, 77(1), 7–9.
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134(1), 9–21.
- Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, 80(2), 312–325.
- Fischer, M., & Shaki, S. (2018). Number concepts: Abstract and embodied. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752), 20170125.
- Fitzgerald, J. K., Freedman, D. J., Fanini, A., Bennur, S., Gold, J. I., & Assad, J. A. (2013). Biased associative representations in parietal cortex. *Neuron*, 77(1), 180–191.

- Ganguli, S., Bisley, J. W., Roitman, J. D., Shadlen, M. N., Goldberg, M. E., & Miller, K. D. (2008). One-dimensional dynamics of attention and decision making in LIP. *Neuron*, *58*(1), 15–25.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155–170.
- Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, *34*(5), 752–775.
- Kanayet, F. J., Opfer, J. E., & Cunningham, W. A. (2014). The value of numbers in economic rewards. *Psychological Science*, *25*(8), 1534–45.
- Kemp, C., & Tenenbaum, J. B. (2008). The discovery of structural form. *Proceedings of the National Academy of Sciences*, *105*(31), 10687–10692.
- Kemp, C., Tenenbaum, J. B., Niyogi, S., & Griffiths, T. L. (2010). A probabilistic model of theory formation. *Cognition*, *114*(2), 165–196.
- King, J. R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: The temporal generalization method. *Trends in Cognitive Sciences*, *18*(4), 203–210.
- Kleiner, M., Brainard, D., & Pelli, D. (2007). "What's new in Psychtoolbox-3?" *Perception*, *36*(14).
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: Integrating cognition, computation, and the brain. *Trends in Cognitive Sciences*, *17*(8), 401–412.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, *40*.
- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, *22*(6), 1027–1038.

- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164(1), 177–190.
- McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: Connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences*, 14(8), 348–356.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518, 529–533.
- Nili, H., Walther, A., Alink, A., & Kriegeskorte, N. (2016). Inferring exemplar discriminability in brain representations. *bioRxiv*.
- Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., & Kriegeskorte, N. (2014). A toolbox for representational similarity analysis. *PLoS Computational Biology*, 10(4).
- O’Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, 15(12), 1–7.
- O’Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, 38(2), 329–337.
- Parkinson, C., Liu, S., & Wheatley, T. (2014). A common cortical metric for spatial, temporal, and social distance. *Journal of Neuroscience*, 34(5), 1979–1987.
- Pisauro, M. A., Fouragnan, E., Retzler, C., & Philiastides, M. G. (2017). Neural correlates of evidence accumulation during value-based decisions revealed via simultaneous EEG-fMRI. *Nature Communications*, 8, 15808.

- Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*, 233–238.
- Schley, D. R., & Peters, E. (2014). Assessing “economic value”: Symbolic-number mappings predict risky and riskless valuations. *Psychological Science*, *25*(3), 753–761.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599.
- Spitzer, B., Waschke, L., & Summerfield, C. (2017). Selective overweighting of larger magnitudes during noisy numerical comparison. *Nature Human Behaviour*, *1*, 1–8.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). Cambridge, MA: MIT press.
- Teichmann, A. L., Grootswagers, T., Carlson, T., & Rich, A. N. (2018). Decoding digits and dice with Magnetoencephalography : Evidence for a shared representation of magnitude. *bioRxiv*.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, *331*(6022), 1279–1285.
- Tervo, D. G. R., Tenenbaum, J. B., & Gershman, S. J. (2016). Toward the neural implementation of structure learning. *Current Opinion in Neurobiology*, *37*, 99–105.
- Twomey, D. M., Murphy, P. R., Kelly, S. P., & Connell, R. G. O. (2015). The classic P300 encodes a build-to-threshold decision variable. *European Journal of Neuroscience*, *42*, 1636–1643.
- Walsh, V. (2003). A theory of magnitude: Common cortical metrics of time, space and quantity. *Trends in Cognitive Sciences*, *7*(11), 483–488.

- Wang, J., Narain, D., Hosseini, E. A., & Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. *Nature Neuroscience*, *21*(1), 102–112.

Chapter 3

Aligned neural geometry between numerical contexts in humans

3.1 Introduction

Humans can think and reason in ways that abstract over the physical properties of the world (Lake, Ullman, Tenenbaum, & Gershman, 2017; Markman & Gentner, 2001). For example, you might know that a cheetah and a jumbo jet can both move “fast” even though felines and aeroplanes belong in different categories and do not look alike. A prevailing view is that this sort of abstract knowledge is supported by neural codes that represent relations among stimuli but are invariant over their physical properties (Gentner, 2010; Summerfield, Luyckx, & Sheahan, 2020; Tenenbaum, Kemp, Griffiths, & Goodman, 2011; Tervo, Tenenbaum, & Gershman, 2016). For example, the concepts of “fast” or “slow” describe how an object moves with respect to the background but do not dictate what that object or background might

be. Encoding relational structure in this way might help agents generalise (or transfer) knowledge from one context to another (Behrens et al., 2018). For example, a decoder trained to discriminate cheetahs from snails might readily be applied to classifying other stimuli as “fast” or “slow”, such as jumbo jets and Zeppelins. Recently, this view has been supported by reports of shared neural geometry among stimuli that bear equivalent relations in different contexts (Baram, Muller, Nili, Garvert, & Behrens, 2019; Bernardi et al., 2019; Luyckx, Nili, Spitzer, & Summerfield, 2019; Remington, Narain, Hosseini, & Jazayeri, 2018).

However, one challenge for this process of structure learning and transfer is that relational patterns in sensory data may be misaligned or rescaled between contexts. We might consider both a record-breaking sprinter and a marathon champion to be “fast” runners, but this does not mean that we expect the latter to run one hundred metres in less than ten seconds, or the former to complete a marathon in under two hours. For generalisation to be effective, we need to form a concept of “fast” that is not tied explicitly to the stopwatch, but that encodes relative information about running speed in each respective context. Fortunately, the brain is prone to normalize incoming sensory signals, with neural activity expressed relative to the local average response (Carandini & Heeger, 2012). Explanations for this ubiquitous computational motif have tended to focus on the way normalization promotes efficient coding by dynamically adjusting the sensitivity of a neural population according to the background (Louie & Glimcher, 2012). We propose, however, that normalization may also serve to realign neural codes in a way that promotes transfer of abstract information across contexts.

To test this hypothesis, we measured brain signals whilst humans viewed streams of symbolic numbers that were sampled (in each block) from one

of three overlapping ranges (or contexts). We chose these stimuli because humans neural patterns for symbolic number are highly structured in human EEG signals, with Arabic digits that convey similar cardinality eliciting correlated patterns of brain activity (Spitzer, Waschke, & Summerfield, 2017). Moreover, there is evidence that this neural code facilitates generalisation of relational information (or structure alignment) (Gentner, 2010) to other numeric (Teichmann, Grootswagers, Carlson, & Rich, 2018) or non-numeric (Luyckx et al., 2019) domains. Using multivariate analyses combined with dimensionality reduction to visualise the neural geometry of the magnitude representation, we observed that neural signals were structured onto 3 equidistant parallel number lines, one for each context. Number lines for the wider range were compressed so that they aligned with the other lines, allowing the brain to signal “more” or “less” in a way that was independent of specific number. In other words, humans appear to learn to use normalization to align neural representations, permitting transfer of abstract knowledge between contexts. We note that these results are preliminary and further analyses need to be conducted to acquire a more complete understanding of the data.

3.2 Methods

3.2.1 Participants

Thirty-nine participants were recruited for the experiment through the participant recruitment system at the Department of Experimental Psychology at the University of Oxford. One participant was omitted from the analyses due to technical difficulties with the EEG equipment. All analyses were performed on the remaining 38 participants (N female = 15, N right-handed = 34, age = $27.11 \pm \text{SD } 6.13$). Participants were required to have normal

or corrected-to-normal vision, with no history of neurological or psychiatric illness. Participants were compensated for their time at a rate of £10 per hour, with an additional maximum bonus of £5 determined by their performance in the task in which they achieved lowest accuracy. The reward contingency was added to ensure participants would pay equal attention to both tasks. Informed consent was given before the start of the experiment. The study was approved by the Medical Science Inter-Divisional Research Ethics Committee (R49432/RE001).

3.2.2 Experimental procedure

Stimuli were created and presented using the Psychophysics Toolbox-3 (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007) for Matlab (MathWorks) and additional custom scripts. The tasks were presented on a 20-inch screen with a resolution of 1600 x 900, at a refresh rate of 60 Hz and on a grey background. Viewing distance was fixed at approximately 62 cm. The ‘up’ and ‘down’ arrow keys on a standard QWERTY keyboard were used as response keys for the numerical comparison task and the space bar for the number matching task.

The experiment was a rapid serial visual presentation paradigm where participants constantly had to alternate between two tasks: a numerical comparison task and a number matching task (Figure 3.1). Participants viewed a sequence of 120 two-digit numbers per block. Coloured numbers were relevant for the numerical comparison task, while white numbers were used only in the number matching task. When participants saw a coloured number (primary target), they had to compare its magnitude to the last coloured number they had seen in the sequence, responding ‘up’ when it was greater and ‘down’ when it was smaller using the arrow keys with their




Figure 3.1. Task design. Participants viewed sequences of 120 rapidly presented two-digit numbers. When a coloured number appeared, they had to compare its magnitude to the last coloured number they had seen, indicated by pressing the 'up' or 'down' arrow key. After a coloured number, 2 to 4 white numbers were presented. Participants were asked to press space bar when a white number matched the last coloured number. Coloured numbers could be drawn from one of three ranges: Low context (25 - 35), High context (30 - 40) or Full context (25 - 40). Which number range was used in a block was determined randomly, with 8 blocks per context.

right hand. Every coloured number was followed by 2 to 4 white 'filler' numbers. Participants then had to press space bar with their left hand when a subsequent white number was identical to the last coloured number (secondary target).

A block started with the presentation of a central fixation cross for 1 second followed by two placeholder hashtag signs for 500 ms. After an interstimulus interval (ISI) of 1 second, the first coloured number was presented centrally. Every stimulus remained at the centre of the screen for 500 ms with a fixed ISI of 1 second after a primary target and a variable ISI between 750 and 1250

ms after a filler number. Participants were able to respond from stimulus onset until 200 ms before the onset of the next stimulus, as to avoid feedback signals overlapping with the presentation of the subsequent stimulus. If a response key was pressed, participants received auditory feedback for 150 ms. Correct responses were followed by a high-pitch tone and any error resulted in a low-pitch tone. If participants failed to respond within the appropriate time window, a buzz sound was played for 150 ms. At the end of a block, participants were informed about their accuracy on both tasks.

The numerical comparison task was split into three contexts, defined by their specific numerical range. In the Low range context, numbers were uniformly drawn between 25 and 35, in the High range context between 30 and 40 and in the Full range context between 25 and 40. The white filler numbers always spanned the entire range between 25 and 40. Each context was associated with a unique colour for the coloured numbers (blue: $\text{RGB} = [86, 180, 233]$; orange: $\text{RGB} = [230, 159, 0]$; purple: $\text{RGB} = [230, 120, 220]$) and the colour-to-context mapping was randomised between participants. The experiment contained 8 blocks per context, resulting in a total of 24 blocks. The block order was randomised in order to reduce temporal similarity between numbers from the same context. A block contained 29 primary targets (excluding the first coloured number of the sequence), with a probability of 12% that at least one of the subsequent white numbers was a secondary target. To prevent potential task switching costs, the white filler number immediately after a primary target was never a secondary target, but participants were not made aware of this feature. At the start of the experiments, participants first completed 3 practice blocks of 144 trials each, one for each of the three contexts. These blocks did not count towards their final performance measures.

3.2.3 Behavioural analyses

All analyses focused on the primary targets. Trials without a response or with response times more than 2.5 SD away from the participant's mean $1/RT$ were also excluded from analysis. First, we tested whether participant's choices were driven not only by the task-relevant decision value but also by the distance to the context mean. Decision value was defined as the absolute difference between the current and previous target. Distance to context mean was a single summary regressor of the difference between context distance of the current target and context distance of the previous target, since both regressors would have opposing signs in the regression. All regressors were z-scored to assess the relative contribution of each regressor. The logistic regression model was estimated for each participant separately and significance of the beta coefficients was tested at the group level (two-tailed t-test, $P < 0.05$). The same regression analysis was repeated, replacing the distance to context mean by a regressor coding for the distance to the global mean ($= 32.5$). Finally, the negative deviances obtained for each model allowed us to compare models using Bayesian model comparison (Stephan, Penny, Daunizeau, Moran, & Friston, 2009). One participant was excluded from the analyses with accuracy as the dependent variable, since they achieved perfect accuracy on the numerical comparison task, which rendered the regression model non-identifiable. Significance tests are thus reported for the 37 remaining participants for analyses on accuracy. To test for contributions in reaction time, the same pipeline was repeated using a general linear model with $1/RT$ as the dependent variable.

To assess normalization in behavioural measurements, we compared two regression models containing a single regressor: the decision value, calculated after normalizing all numbers between 0 and 1 either within each context

(relative scale) or within the full range of the experiment (absolute scale). Negative deviances were entered into a Bayesian model selection algorithm to determine the best fitting model. A logistic regression model was used when predicting accuracy and a general linear model for reaction time.

3.2.4 EEG acquisition

The EEG signal was recorded using 61 Ag/AgCl sintered surface electrodes (EasyCap, Herrsching, German), a NeuroScan SynAmps RT amplifier, and Curry 7 software (Compumedics NeuroScan, Charlotte, NC). Electrodes were placed according to the extended international 10-20 system, with the right mastoid as recording reference and channel AFz as ground. Additional bipolar electrooculography (EOG) was recorded, with two electrodes placed on either temple for recording horizontal EOG and two electrodes above and below the right eye for vertical EOG. All data was recorded at 1 kHz and low-pass filtered online at 200 Hz. All impedances were kept below 10-15 k Ω during the experiment.

3.2.5 EEG pre-processing

The data were pre-processed using functions from the EEGLAB toolbox (Delorme & Makeig, 2004) for Matlab and custom scripts. First the data were down-sampled to 250 Hz, low-pass filtered at 40 Hz and then high-pass filtered at 0.5 Hz. The continuous recording was visually screened for excessively noisy channels and these were interpolated by the weighted average of the surrounding electrodes. The data was then offline re-referenced to average reference. Epochs were extracted from 250 ms before to 1000 ms after stimulus onset. Epochs were baselined relative to the full pre-fixation time window. Epochs containing atypical noise (such as muscle activity) were

rejected after visual inspection. We then performed Independent Component Analysis (ICA) and removed components related to eye blink activity and other artefacts (manually selected for each participant).

3.2.6 EEG: Univariate analyses

A regression-based approach allowed us to disentangle the influence of various variables on the amplitude changes in neural signals during the numerical comparison task. Before running the regression, we used Principal Component Analysis (PCA) on the pre-processed data of each participant to reduce noise in the signal. All epochs with primary targets were stacked over all electrodes and trials, creating a feature matrix of (trials*electrodes) by time points. The first 43 principal components (PC) were retained, which on average explained 90% of the data. The EEG data was then reconstructed by multiplying the PC's with the estimated loadings and the reconstructed data was subsequently used as the dependent variable in our linear regression model.

The regression model contained 4 regressors of interest: (1) numerical magnitude of the current primary target, (2) the absolute difference of the current primary target to the previous target, (3) absolute difference between the current target and the mean of the current context, and (4) absolute difference between the previous target and the mean of the current context. We included two nuisance regressors: (5) visual similarity between the current primary target and the previous target and (6) reaction time (RT) on the current trial. Visual similarity was estimated as the correlation distance between black-and-white pixel images of two numbers as they appeared on screen. Trials with no response or RTs beyond 2.5 SD of the inverse RT were excluded from analysis. All regressors were z-scored before estimating the beta coefficients for each time point and electrode.

3.2.7 EEG: Representational Similarity Analysis (RSA)

To understand the representational structure contained in the multivariate neural signal during the numerical comparison task, we constructed neural Representational Dissimilarity Matrices (RDMs) at each time point. First, all data was z-scored over all trials per electrode and time point. Condition-specific activations were estimated at each time point and electrode through a regression model with dummy coding for every number per context (in total 38: 11 for Low, 11 for High and 16 for Full), where the beta coefficients reflected the average activations per condition. The residuals of the regression were used to estimate the covariance matrix, which was subsequently used to increase the signal-to-noise ratio (SNR) by noise normalizing (whitening) the averaged EEG data through multiplication with the negative half inverse covariance matrix ($\Sigma^{-0.5}$) at each time point (Walther et al., 2016). Finally we calculated the Pearson correlation distance between each condition, resulting in a 38 x 38 RDM at each time point. To exclude the possibility that more observations in certain contexts were biasing the dissimilarity measures, we constructed RDMs iteratively with subsampled data that equated the frequency of observations in each cell of the RDM. We repeated the RDM construction 100 times, randomly selecting N observations per cell at each iteration, with N determined by the minimum number of observations in any condition for a particular participant. The final neural RDM then consisted of the averaged RDM over all 100 iterations.

We constructed 3 candidate model RDMs that represented different potential structures in the neural signal: context, magnitude and visual similarity. The context model RDM assumed complete similarity ($= 0$) between numbers coming from the same range context and complete dissimilarity ($= 1$) for numbers coming from different ranges. The magnitude model RDM assumed

a relative magnitude code, placing the numbers within a range between 0 and 1 and subsequently calculating the Euclidean distance between numbers from all ranges. The model RDM representing the visual similarity between two numbers was calculated using correlation distance between the vectorized black-and-white pixels of the numbers as they appeared on screen. All model RDMs were z-scored before entering in the linear regression to assess their relative contribution. The neural RDMs were further smoothed with a 60 ms sliding time window. Beta coefficients were estimated at each time point and for each participant separately. Statistics are reported at the group level. Significant clusters were identified using cluster-corrected non-parametric permutation tests (iterations = 5000, cluster-defining and cluster-level thresholds at $P < 0.05$, unless stated otherwise) (Maris & Oostenveld, 2007).

For a supplementary control analysis, two additional subject-specific models were added to the regression: a colour model RDM and a reaction time model RDM. For the model RDM representing the colour space, RGB values were first transformed into CIE 1976 L*a*b* space to more accurately reflect human perception of colour differences. The dissimilarity between colours was then indexed through ΔE^* , a measure of the colour difference between two colours in CIELAB space. The RT RDM was constructed by calculating the average RT to a number independent of the preceding number. RTs were cut off based on the 2.5 standard deviation from the inverse RT to control for outliers before averaging.

3.2.8 Line configuration model

We created a function that could simulate different three-dimensional configurations of the three number lines, ranging from fully orthogonal to parallel (see Figure B.1). Each number line was defined by two angle values, θ and φ .

The length of each line was set at 1, centred at the origin. The x-, y-, and z-coordinates for the end points of each line was then calculated as follows:

$$\begin{aligned} x &= .5 \times \sin(\theta) \times \cos(\varphi) \\ y &= .5 \times \sin(\theta) \times \sin(\varphi) \\ z &= .5 \times \cos(\theta) \end{aligned} \tag{3.1}$$

Numbers were equally spaced on a line interpolated between $(-x, -y, -z)$ and (x, y, z) stored in $\mathbf{X}$, $\mathbf{Y}$ and $\mathbf{Z}$. To allow for separation between the three lines, the centre of each interpolated line was positioned on a corner of an equilateral triangle in the XY-plane with the origin at its centre and the distance from the origin C could freely vary between $[0, 2]$. All x- and y-coordinates were adjusted accordingly:

$$\begin{aligned} \mathbf{Y}_{low} &= \mathbf{Y}_{low} + \frac{C}{\sqrt{3}} \\ \mathbf{X}_{high} &= \mathbf{X}_{high} - \frac{C}{2}, \mathbf{Y}_{high} = \mathbf{Y}_{high} - \frac{C \times \tan(\frac{\pi}{6})}{2} \\ \mathbf{X}_{full} &= \mathbf{X}_{full} + \frac{C}{2}, \mathbf{Y}_{full} = \mathbf{Y}_{full} - \frac{C \times \tan(\frac{\pi}{6})}{2} \end{aligned} \tag{3.2}$$

The final model thus contained 7 free parameters (6 angles and 1 distance parameter). Finally, to fit the 3D line configurations to the neural data, we calculated the Euclidean pairwise distance between the numbers on all lines to create a model RDM. The neural data were constructed by determining for each participant a time window of significant magnitude representation through leave-one-out estimation, repeating the RSA without the participant of interest and extracting the time window in which magnitude was significantly represented ($P_{cluster} < 0.05$). This allowed us to obtain an independent

time window estimate for each participant. Both the line configuration RDM and the averaged neural RDM were z-scored and the best fitting model was found through least squares, using a scatter-based algorithm with 50 random starting points.

The estimated angle parameters allowed us to compare the three vectors underlying the number lines and thus compare how parallel they were. We calculated the cosine similarity ($\cos(\theta)$) between all pairs of vectors. If $\cos(\theta) = 0$, then the two vectors are orthogonal; if $\cos(\theta) = 1$, the vectors are parallel and if $\cos(\theta) = -1$ the vectors point in opposite directions.

3.2.9 Line normalization model

After establishing that the three number lines were in fact on average parallel, we assessed whether these lines were represented along a single number line (absolute scale) or whether numbers were squashed within a fixed range (relative scale). For each participant we tested 3 models: relative scale, absolute scale and free varying scale. The model contained 6 parameters: 3 parameters determining the length of each number line and 3 parameters for the roll, pitch and yaw of all three lines. First, the three number lines were initialized as points in a three-dimensional space along the z-axis, denoted as M . In the absolute model the spacing between numbers in the full range context was constrained to be 16/11, while spacing was equal to 1 for all contexts in the relative model or a free parameter in the free model. Next, the lines were rotated through multiplication with the rotation matrix $R(\alpha, \beta, \gamma)$ to obtain the rotated configuration L that best described the angle of the three lines in the neural data, where α , β and γ determined the yaw, pitch and roll respectively. Finally, the lines were separated by the empirical distances obtained through MDS, i.e. the average neural distance between any two

contexts. We again constructed model RDMs from the three-dimensional configuration and estimated the best fitting model through least squares using a scatter-based algorithm with 50 random starting points ($P_{cluster} < 0.05$).

As an initial indication of normalization, we plotted the MDS scores on the third dimension for the group-averaged RDM between 500 - 800 ms against a number space between 0 and 1 for each context (Figure 3.4D). Spearman rank correlations in each context were used to indicate the presence of a magnitude code in the third dimension. Next we wanted to test normalization by using the slope estimated in one context and evaluating its fit in the other contexts. The abscissa was either spaced according to a relative scale (all squashed between 0 and 1) or an absolute scale. If number lines were normalized within contexts, estimated slopes should be similar across contexts. We calculated the cross-validation error by using the slope estimated in one context to predict scores in the other two contexts, repeated for relative and absolute scaling. The reported error was the sum of squared errors over all cross-validation iterations.

Finally, to allow statistical tests of the parameters estimates, we fitted the normalization models to individual participant's neural RDMs (averaged over each individual's cross-validated time window). We compared the loss for a relative and absolute model using Bayesian model selection. For the freely varying model we ran a repeated-measures ANOVA on the the estimated line lengths in each context to test for any consistent differences in line length.

3.2.10 Cross-generalisation analysis

One possible explanation for the separation between contexts in the neural data is due to a similarity in colour and/or time. We therefore turned to a decoding approach with the filler numbers as an independent measure of

context. Filler numbers (excluding secondary targets that required a response) always covered the full range of numbers (25 - 40) and were all coloured in white, thus making them ideal as independent training data to a classifier by context. We used Linear Discriminant Analysis (LDA), a classifying algorithm that attempts to find a linear combination of features (in this case neural signals over electrodes) that separates contexts. To control for similarity over time and improve decoding accuracy (Grootswagers, Wardle, & Carlson, 2017), training and test data were constructed by randomly combining epochs from different blocks of the same context, resulting in 24 training and 24 test data points (8 per context, equal to the number of blocks in the experiment). Individual epochs were smoothed with a 60 ms sliding time window before averaging. The averaged training and test epochs were noise normalized to improve SNR. The classifying algorithm was then trained at each time point of the training data and evaluated at each time point of the test epochs. Decoding accuracy was determined by selecting the context with the highest classification probability and comparing the model generated answer to the true answer. The cluster-correction procedure described above was used for the two-dimensional data to identify above-chance (> 0.33) clusters at the group level ($P_{cluster} = 0.005$) (Maris & Oostenveld, 2007). We note that the decoding analysis is preliminary and will be revised in the future to better exclude the possibility of temporal similarity.

3.3 Results

3.3.1 Context-driven processing

Healthy human adults ($n = 38$) performed a numerical comparison task whilst scalp EEG signals were recorded (Figure 3.1). The task required them to

indicate whether a target two-digit Arabic number was greater or smaller in magnitude than the previous target number. Target stimuli occurred in each of three contexts, which spanned digit ranges 25-35 (Low), 30-40 (High) and 25-40 (Full) in distinct blocks. Targets were presented in a colour that signalled the context in which they occurred (counterbalanced across participants) and were interspersed with 2 - 4 “filler” numbers (all white) that were drawn from across the full range. To maintain attention to the fillers, participants performed a secondary task that did not require magnitude comparison, namely they pressed the space bar when a filler exactly matched the previous target ($\sim 12\%$ of fillers).

We focus on the primary (magnitude comparison) task which participants performed with a mean accuracy of 94 ($\pm$ SD 4)% and response times that averaged 675 ($\pm$ SD 63) ms. Confronted with each target, there are two (potentially redundant) sources of information that can be used to guide responding: the identity of the previous target, and the range of numbers in the current context. Intuitively, a participant who forgets the previous target number but responds “larger” to relatively higher numbers will perform above chance. Behavioural analyses revealed that participants used both of these sources of information in making decisions. Logistic regression revealed that accuracy was explained by both the disparity between current and past target number (target distance; two-sided t-test, $t_{36} = 7.60, P < 0.001$) and the distance between the target number and the median number in the context (context distance; $t_{36} = 5.56, P < 0.001$) (Figure 3.2A). Target distance ($t_{37} = 5.79, P < 0.001$) but not context distance ($t_{37} = 1.01, P = 0.32$) significantly explained variance in reaction time. When we compared models in which the global context (median across all possible numbers) and local context (median across the current context only) were used as an auxiliary




Figure 3.2. Context-dependent processing. **(A)** A forgetful agent would achieve higher accuracy if it used the mean of the current context as a second determinant of choice. Logistic regression on accuracy revealed that participant’s choices in the number comparison task were not only explained by target distance (difference between current and previous primary target), but also by distance to the mean of the context. Error bars denote the 95%-confidence interval **(B-D)** A regression analysis of scalp activations further confirmed that a centro-parietal region encoded the decision value (DV) of target distance, while right prefrontal electrodes encoded the distance to the mean of the context, both for the current and previous target. Significant clusters were defined using FDR correction at 0.05 for visualization purposes.

source of information, the latter explained human choices better (Bayesian model selection, protected exceedance probability (pxp) = 0.92) (Stephan et al., 2009).

Next, we examined the EEG data using a univariate approach, because previous studies have shown positive potentials over the parietal cortex covary inversely with the uncertainty prompted by a decision (Kelly & O’Connell, 2015; O’Connell, Dockree, & Kelly, 2012; Wyart, de Gardelle, Scholl, & Summerfield, 2012; Wyart, Myers, & Summerfield, 2015). We thus regressed the target and context distance predictors onto single-trial EEG data, mapping the positive effects and visualising the beta time series over

parietal electrodes (Figure 3.2B-D). Consistent with previous findings, the (unsigned) difference between the two target numbers was positively encoded over posterior parietal electrodes from $\sim 250 - 300$ ms post-stimulus onset (non-parametric cluster correction, $P_{cluster} < 0.05$) (Maris & Oostenveld, 2007), but we also observed a consistent representation of (local) context distance over right prefrontal electrodes for both the previous ($\sim 150 - 300$ ms) and the current number ($\sim 250 - 300$ ms; $P_{cluster} < 0.05$). Taken together the behavioural and univariate results show that humans learned the range of numbers in each context and used it to make decisions.

3.3.2 Multivariate representation of context and magnitude

We now turn to our main questions of interest, which concerned the multivariate neural representation of magnitude and context. We constructed 3 representational dissimilarity matrix (RDM) models that encoded (a) magnitude differences among the numbers and (b) differences among contexts, as well as (c) a nuisance RDM that encoded physical (pixel-wise) similarity in the Arabic digits. We regressed a neural RDM obtained from scalp EEG patterns at each time point post-target against the three model RDMs. Whilst early time points were dominated by visual similarity (Figure 3.3A), from about 400 ms post-stimulus the neural patterns diverged with differences in both context (Figure 3.3B) and magnitude (Figure 3.3C) (all $P_{cluster} < 0.05$). A supplementary control analysis suggested that these results were not entirely driven by motor components or colour similarity within conditions, after adding 2 nuisance model RDMs for (a) reaction time and (b) colour similarity to the original regression model (Figure B.1).

Because human colour perception is hard to approximate, we wanted to obtain an independent measure of context separation. We used the filler




Figure 3.3. Context and magnitude representation in multivariate signals. (A-C) Neural representational similarity matrices (RDMs) at each time point were regressed against three model RDMs (see insets) for (A) visual similarity, (B) context, and (C) magnitude. Early multivariate representations resembled more visual similarity structure, while a magnitude representation arose from around 400 ms. Shaded area represents SEM. Lines below show significant clusters after cluster-correction ($P_{cluster} < 0.05$). (D) Cross-generalisation from filler numbers to primary targets. We observed above-chance decoding (> 0.33) of context between 250 to 750 ms, indicating that context separation was not purely driven by temporal or visual features of the task. (E) Early univariate signals in right posterior electrodes encoded the numerical magnitude of the current target. Significant clusters were defined using FDR correction at 0.05 for visualization purposes.

numbers from the number matching task as an independent training set to a linear classification algorithm (Linear Discriminant Analysis, LDA), since fillers were all white, spanned the full range of numbers and required a different motor response to the numerical comparison task. The classifier was trained at a single time point from the filler epoch and tested at each time point of the primary target epoch. We found above-chance decoding (> 0.33) between $\sim 250 - 750$ ms after stimulus onset for the primary target ($P_{cluster} < 0.005$), suggesting the separation by context was not *solely* driven by external features of the task (Figure 3.3D).

Finally, because neural RDMs were calculated using correlation distance, which corrects for mean differences between conditions, the observed effects should not have been driven by univariate effects. We do note that we detected an earlier effect of magnitude in univariate signals between 100 and 300 ms post-stimulus localised in right posterior electrodes (Figure 3.3E) and these results could not be explained by the visual similarity in shared decades between targets given the additional nuisance regressor.

3.3.3 Aligned geometry in neural signals

Next, we used multidimensional scaling (a dimensionality reduction technique) to visualise relations among the neural codes associated with each number in each context (Figure 3.4 A). In Figure 3.4 B we show the projection of each number in each context into a space spanned by the first three dimensions of the neural similarity data. As can be seen, the numbers spread out along parallel lines in the third dimension, whereas the three contexts themselves were separated in the first two dimensions, lying approximately at the apices of an equilateral triangle (left panel). To more formally test the assertion that lines were in fact parallel, we developed a model that could independently vary

the angle of three lines in a three-dimensional space (see Methods), similar to the MDS configuration. Line configurations were translated into RDMs that could subsequently be used to fit to the neural data using least squares and a scatter-based search algorithm with 50 random initializations. After obtaining the best fitting angles for each participant, we calculated the pairwise cosine distance between any two contexts. A cosine distance of 1 implies that the lines are pointing in the same direction, while 0 indicates orthogonal lines and -1 opposing directions. We found that all cosine distances were significantly greater than 0 for any pairwise comparison between fitted contexts at the group level (one-tailed Wilcoxon signed rank test, all $P < 0.008$) (Figure 3.4 C).

Visual inspection suggested that the number representations within each line were compressed and/or recentred so that the third dimension expresses a relative quantity (i.e. “more” or “less” within each context) rather than just cardinality. To verify this, we plotted the stimuli ordered by magnitude against their average value on the 3rd dimension, separately for each context, and fit a linear trend (Figure 3.4 D). First of all, the slopes were positive in each context, implying the numbers are ordered along each line from low to high, and rank correlations between number and dimension 3 scores were all significant (all $P < 0.002$). Importantly however, when the abscissa was scaled from the minimum to maximum value in each context, the linear slopes were approximately equal, implying that number representations were normalized. We tested this more formally by fitting a linear trend to the third dimension of neural data of each context and using it to decode the relative or the absolute number in the held-out contexts. The cross-validation error over all ranges for the absolute model ($SSE = 0.92$) was substantially higher than for the relative model ($SSE = 0.31$). This confirms what is implied by




Figure 3.4. Aligned geometry of number lines. **(A)** Using multidimensional scaling (MDS), we visualized the first three principal components of the neural data at the group level, averaged between around 500 - 800 ms. Numbers lay on parallel lines separated by context. **(B)** We projected the dimensions against each other to better understand the low-dimensional configuration. The first two components showed a clear separation by context (left panel), while the scores on the third dimension suggested a common axis of magnitude. **(C)** Using a model that could assume a parallel 3D configuration we were able to estimate the angle of the number lines for each participant. Estimated cosine distances on average were close to 1, indicating parallel lines. Density plots used a kernel with bandwidth of 0.02. **(D)** Line fits to MDS scores of the third dimension showed equal slopes between contexts when the abscissa was equally spaced between minimum and maximum number (relative code).

the model-based analysis above: that the representation numbers 25 - 40 in the full context is “squashed” into the same neural space as numbers 25 - 35 and 30 - 40 in the low and high contexts respectively. In other words, the brain is engaging in a form of normalization in the low-dimensional neural space in which symbolic number is represented.

Finally, we assessed the statistical reliability of these effects at the group level by parametrically generating model RDMS using the line normalization model and fitting these to the neural data for individual participants. We compared a relative model in which the elongation of the line was equated in all three contexts, irrespective of the differing ranges of numbers to an absolute model where line length was proportional to the range of numbers that occurred in that context. Using Bayesian Model Selection at the group level to compare the relative and absolute models head to head, we found that protected exceedance probability favoured the relative model, consistent with the observations above ($pxp = 1$; expected frequencies for relative model were 0.99 vs. 0.01). However, when we fit a free model in which the line length was unconstrained in each context and compared the best-fitting line lengths we found the full context to have a consistently longer length ($M_F = 0.95 \pm 0.02$) than the other two contexts ($M_L = 0.89 \pm 0.05$, $M_H = 0.90 \pm 0.04$) (repeated-measures ANOVA, $F_{(2,74)} = 18.81, P < 0.001$). Estimated line lengths were longer for the Low and High context and closer to the Full context than would be expected from a truly absolute scale. These results need further investigation. Importantly the normalization of number lines was also reflected in behaviour. When we compared two logistic regression models that explained variance using the absolute difference between current and previous target on either an absolute or relative scale, we found that Bayesian model selection favoured a relative model over an absolute model, both for

accuracy ($pxp_{rel} = 0.64$) and reaction time ($pxp_{rel} = 0.57$).

3.4 Discussion

The neural results described here suggest three principles of conceptual encoding. Firstly, as previously described (Luyckx et al., 2019; Spitzer et al., 2017; Teichmann et al., 2018), multivariate neural signals recorded at the scalp encode the magnitude that is signalled by symbolic numbers. Secondly, numbers presented in different contexts are segregated by a low-dimensional offset, such that they are arranged into parallel number lines. This mirrors findings from work in non-human primates (Bernardi et al., 2019; Remington et al., 2018). In our data, the offset was approximately 2-dimensional, such that numbers in each of the contexts could be visualised in three equidistant parallel columns. Thirdly, magnitude is likely coded on a relative scale within each number line, so that the ends of the line signalled “more” or “less” rather than a specific numerosity.

The neural coding patterns described here seem to meet the computational desiderata introduced earlier in this article: relative coding of magnitude in parallel neural lines might allow observers to read out “more” and “less” in different contexts, permitting generalisation of magnitude information but avoiding interference. While it was indeed the case that magnitude was represented along a single axis with similar spread, which should facilitate the read-out of numerical magnitude independent of context, our findings are preliminary and apply only to ‘near transfer’: the task was identical for each context and always applied to numbers. Future work should be carried out that tests the case where transfer is needed between sets of stimuli that share the same underlying structure, but require different tasks or contain items with different features. These experiments would also be able to study

how well our results generalize to other structured input beyond numerical magnitude, since our adult participants were very familiar with the use of numbers. Luyckx et al. (2019) found that neural signals for value aligned with numerical magnitude, a first indication that structure can translate to newly learned information that does not share any physical properties (see Chapter 2). However their experiment was not built to test the normalization hypothesis and required different tasks for each stimulus set.

Several analyses are in the pipeline that should help us understand the data better. One of the next steps is to integrate the line configuration and line normalization model into one comprehensive model that can test the aligned geometry of the number lines in combination with the normalization of the number lines. We have also started work on recurrent neural networks (RNN) that can mirror the task structure, so that numbers are kept in memory while a secondary task is being performed. Preliminary results suggest that context separation is present only when context is explicitly signalled (as was the case in our task with distinct colour mapping), that the network uses the context mean as a guide when the network was lesioned on a portion of input during training (mimicking a forgetful learner) and normalizes number lines when contexts are grouped in blocks. These findings would be in line with the idea that participants use the context when they are forgetful. More work is needed here and we hope to present these results soon.

References

- Baram, A. B., Muller, T. H., Nili, H., Garvert, M., & Behrens, T. E. J. (2019). Entorhinal and ventromedial prefrontal cortices abstract and generalise the structure of reinforcement learning problems. *bioRxiv*.
- Behrens, T. E., Muller, T. H., Whittington, J. C., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, 100(2), 490–509.
- Bernardi, S., Benna, M. K., Rigotti, M., Munuera, J., Fusi, S., & Salzman, C. D. (2019). The geometry of abstraction in hippocampus and prefrontal cortex. *bioRxiv*.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, (10), 433–436.
- Carandini, M., & Heeger, D. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1), 51–62.
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134(1), 9–21.
- Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, 34(5), 752–775.
- Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time series neuroimaging data. *Journal of Cognitive Neuroscience*, 29(4), 677–697.
- Kelly, S. P., & O’Connell, R. G. [R. G.]. (2015). The neural processes underlying perceptual decision making in humans: Recent progress and future directions. *Journal of Physiology Paris*, 109(1-3), 27–37.

- Kleiner, M., Brainard, D., & Pelli, D. (2007). "What's new in Psychtoolbox-3?" *Perception*, 36(14).
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.
- Louie, K., & Glimcher, P. W. (2012). Efficient coding and the neural representation of value. *Annals of the New York Academy of Sciences*, 1251(1), 13–32.
- Luyckx, F., Nili, H., Spitzer, B., & Summerfield, C. (2019). Neural structure mapping in human probabilistic reward learning. *eLife*, 366757.
- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164(1), 177–190.
- Markman, A. B., & Gentner, D. (2001). Thinking. *Thinking*, 52(1), 223–247.
- O'Connell, R. G. [Redmond G], Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, 15(12), 1–7.
- Remington, E. D., Narain, D., Hosseini, E. A., & Jazayeri, M. (2018). Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron*, 98(5), 1005–1019.e5.
- Spitzer, B., Waschke, L., & Summerfield, C. (2017). Selective overweighting of larger magnitudes during noisy numerical comparison. *Nature Human Behaviour*, 1, 1–8.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Group Bayesian Model Selection for Group Studies. *NeuroImage*, 46(4), 1004–1017.
- Summerfield, C., Luyckx, F., & Sheahan, H. (2020). Structure learning and the posterior parietal cortex. *Progress in Neurobiology*, 184, 101717.

- Teichmann, A. L., Grootswagers, T., Carlson, T., & Rich, A. N. (2018). Decoding digits and dice with magnetoencephalography: Evidence for a shared representation of magnitude. *Journal of Cognitive Neuroscience*, *30*(7), 999–1010.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, *331*(6022), 1279–1285.
- Tervo, D. G. R., Tenenbaum, J. B., & Gershman, S. J. (2016). Toward the neural implementation of structure learning. *Current Opinion in Neurobiology*, *37*, 99–105.
- Walther, A., Nili, H., Ejaz, N., Alink, A., Kriegeskorte, N., & Diedrichsen, J. (2016). Reliability of dissimilarity measures for multi-voxel pattern analysis. *NeuroImage*, *137*, 188–200.
- Wyart, V., de Gardelle, V., Scholl, J., & Summerfield, C. (2012). Rhythmic fluctuations in evidence accumulation during decision making in the human brain. *Neuron*, *76*(4), 847–858.
- Wyart, V., Myers, N. E., & Summerfield, C. (2015). Neural mechanisms of human perceptual choice under focused and divided attention. *Journal of Neuroscience*, *35*(8), 3485–3498.

Chapter 4

Robust decision-making through normalization and compression under noisy processing

4.1 Introduction

When cognitive scientists describe human decision-making, it is often interpreted in comparison to a normative standard of how one should act. Optimal behaviour, i.e. making choices that lead to the best possible outcome, is then assumed to be achieved if one follows the normative solution. However, when the normative theory is ill-defined, rational choices will appear as suboptimal or irrational. It has been argued that the apparent suboptimality in perceptual and economic decision-making arises from the fact that optimality is defined in terms of Bayes optimality, which does not take into account any biological constraints of the system, such as energy cost or cognitive

resource limitations (Moran & Tsetsos, 2018; Sims, 2006; van den Berg & Ma, 2018; Wei & Stocker, 2015). For example, a large portion of choice variability can be explained by ‘late noise’ or noise in the inferential process (Drugowitsch, Wyart, Devauchelle, & Koechlin, 2016; Wyart & Koechlin, 2016), which is distinct from sensory or encoding noise assumed in Bayesian models (Beck, Ma, Pitkow, Latham, & Pouget, 2012). Consequently, any behaviour diverging from the Bayesian model predictions will be regarded as ‘suboptimal’.

The theory of efficient coding on the other hand postulates that our brain attempts to maximize efficiency of processing information given certain resource limitations (Barlow, 1961; Bhui & Gershman, 2018; Polania, Woodford, & Ruff, 2019; Wei & Stocker, 2015). One way this can be achieved is by discarding redundant or outlying information to obtain an estimate that is robust to noisy integration. Computational and empirical work have demonstrated that while at face value such a strategy might appear to lead to suboptimal choices, it in fact maximizes performance given certain limitations of the system (Bhui & Gershman, 2018; Li, Castañón, Solomon, Vandormael, & Summerfield, 2017; Spitzer, Waschke, & Summerfield, 2017; Summerfield & Tsetsos, 2015; Tsetsos et al., 2016). For example, Spitzer et al. (2017) presented sequences of 10 digits randomly drawn from one of two colour categories, and asked participants to choose the colour category with the highest mean. They found that rather than weighting all numbers equally, as would be suggested under a normative framework, participants overweighted larger numbers in their decisions. Simulations showed that, if one assumes information is corrupted by late noise, an anti-compression of the number line results in better performance. Moreover the level of anti-compression could be extracted from neural representations of the number line. These findings

were replicated by Luyckx, Nili, Spitzer, and Summerfield (2019) (Chapter 2, section 2.3.2) both at the behavioural and neural level, and were shown to be independent of task framing (i.e. choosing the colour category associated with either the highest or lowest average). Tsetsos et al. (2016) in their model fits reported a strong relation between the level of estimated late noise and the level of estimated 'selective gating', a parameter that governed how linear the integration of information was. In other words, participants with a noisier system compensated with stronger non-linearities in their accumulation of evidence (for a replication of these results, see Chapter 5).

An important aspect of efficient coding these studies have so far failed to highlight is that the brain adapts to the range of information in the environment. Originally it was found that neurons can dynamically change their sensitivity to adapt to the average level of stimulation, for example when adapting to the light intensity (Carandini & Heeger, 2012). Sufficient evidence now exists to suggest that a similar motif of divisive normalization exists beyond sensory processing for higher-level concepts such as value-based processing (Louie & Glimcher, 2012). Previous studies have demonstrated that humans can flexibly adapt to different portions of the stimulus space to optimise performance (de Gardelle & Summerfield, 2011; Michael, de Gardelle, Nevado-Holgado, & Summerfield, 2013), but they did not employ a sequential sampling paradigm where late or accumulation noise could play a strong factor. On the other hand the sequential sampling studies mentioned above employed a limited range of single-digit numbers and were thus not equipped either to determine the impact of normalization. Here we capitalize on the non-linearities in information-processing found in those previous studies to distinguish between absolute and relative (normalized) processing of information when using multiple ranges in a sequential sampling

paradigm. Through model simulations and model fitting we show that aside from anti-compression, normalization grants more robust solutions in a system with late noise.

4.2 Methods

4.2.1 Participants

We recruited 24 healthy participants through the participant recruitment system of the Department of Experimental Psychology at the University of Oxford. All participants had normal or corrected-to-normal vision, with no history of neurological or psychiatric illness and should not have taken part in the experiment of Luyckx et al. (2019). Three participants were omitted from the analysis due to chance level performance. All analyses were performed on the remaining 21 participants (N female = 14, N right-handed = 21, age = $25.43 \pm \text{SD } 5.65$). Participants were paid £10 for completing the experiment and earned a monetary bonus of maximum £2.5 based on accuracy. Informed consent was given before the start of the experiment. The study was approved by the Medical Science Inter-Divisional Research Ethics Committee (R49432/RE001).

4.2.2 Experimental procedure

The task design was an extension of the numerical integration task described in Luyckx et al. (2019) (Chapter 2, section 2.2.2, p. 36). The trial structure differed only slightly, with the location of response options on screen (left vs right response box) now randomized from trial-to-trial, instead of fixed per participant. Participants all received the same task framing where they had to indicate which of the two colour categories had the highest average

(Figure 4.1).




Figure 4.1. Task design. Participants performed a numerical integration task where they had to indicate the colour category with the highest average over a sequence of 10 two-digit numbers. Colour assignment to any number was random but balanced within each sequence. Numbers were drawn from either a 'low' or 'high' range, each containing numbers in six equidistant steps of 2. The first number of a range was randomly assigned for each participant (see bottom). The two ranges were further highlighted through distinct colour mappings (orange/blue and purple/green). Aside from 2 blocks per range, participants also completed 4 blocks where sequences from the two ranges were intermixed. Numbers below the frames indicate timing in milliseconds.

The experiment consisted of 32 practice trials (8 Low, 8 High, 16 Mixed) and 400 experimental trials, divided in 8 blocks of 50 trials. There were 2 blocks for the Low and High range each, and 4 blocks of Mixed blocks. This resulted in 100 trials per range and block type combination (100 Low, 100 High, 100 Mixed Low and 100 Mixed High). Each block started with

an indication of the block type ('Low', 'High', 'Mixed'). Block types were randomized across the experiment, except for the first block, which would always either be the Low or High condition, never the Mixed.

We increased the numerical values compared to the original task over two different ranges: 'Low' and 'High' range. We also included a third block type where trials from the two ranges were randomly interleaved ('Mixed'). The low range spanned 6 numbers in steps of 2 with a starting value randomly determined for each participant between 23 to 25. Similarly, the high range could start at 35 up to 37, randomly assigned per participant. The random starting point was added to ensure that the natural inflection point around multiples of 10 (30 or 40) was not biasing our results. The two ranges were indicated by a different colour combination (blue/orange or green/purple) to further emphasize the difference between the two ranges. Colours were assigned to conditions based on participant number (i.e. odd participant numbers: low = o/b, high = g/p, and vice versa for even participant numbers).

4.2.3 Psychometric models and model fitting

Since the experiment now consisted of different ranges, we hypothesized several extensions of the original psychometric model (Chapter 2, section 2.2.10, p. 46) that could explain the data. First, numerical values were either normalized between 0 and 1 within the range of the entire experiment (absolute) or within their respective range (relative). Second, the distortion of the number line (k) was either fixed over all conditions ($n_k = 1$), differed between low and high range trials regardless of block type ($n_k = 2$), differed between block types ($n_k = 3$) or differed between range trials and block types ($n_k = 4$). The remaining three parameters (motor bias B , leak l and late noise s) were assumed to be fixed over all conditions within a participant.

The parameters search space was restricted as follows:

$$B \in \{-\infty, \dots, +\infty\}$$

$$s \in \{0.001, \dots, 8\}$$

$$l \in \{0, \dots, 1\}$$

$$k \in \{e^{-2.5}, \dots, e^{2.5}\}$$

We used 5-fold cross-validation to determine which of the 8 candidate models was most likely to explain the choice data. This approach allowed a fair comparison between models with different numbers of parameters. For each iteration, we estimated best fitting parameter combinations based on 80% of trials, with an equal portion of trials coming from different conditions when $n_k > 1$. Parameter estimates were obtained by minimizing the negative sum of the log-likelihoods using a scatter-based global optimization solver in Matlab. We then predicted choices on the remaining 20% test trials using the parameter estimates from the training set to test the goodness of fit. After leaving out each subset of test trials, log-likelihood estimates from all test trials were then summed and models were compared by calculating the protected exceedance probability (pxp) through a Bayesian Model Selection procedure (Stephan, Penny, Daunizeau, Moran, & Friston, 2009). Final parameter estimates for the best fitting model were then obtained by repeating the same fitting procedure on the full dataset.

4.2.4 Model simulations

We first generated 10 000 new trials, 2500 for each range and block type. Similar to the real experiment, sequences consisted of 10 numbers with

category assignment balanced within each sequence. Trials were replaced when the category means were equal. Because ranges could start at different values, we set the ranges for the simulations to 23 - 33 (Low) and 37 - 47 (High) in equidistant steps of 2, to mimic the case of a participant where the ranges were maximally distant. Numbers were then normalized between 0 and 1, either within their respective range (relative model) or within the full range of the experiment (absolute model). Next, number lines were compressed by exponentiating them with a value of k , as in Equation 2.3 (p. 46). Because of the non-linear transformation, deviations from linear will lead to larger (or smaller) absolute decision values, which in turn will bias the estimated model accuracy. Moreover, from a biological perspective this would suggest that neurons in the limit can infinitely increase their firing rate. We therefore normalized the transformed values with a multiplicative gain factor g :

$$g = \frac{\sum f^k}{\sum f} \quad (4.1)$$

with f representing the feature space, i.e. the unique untransformed numbers, within a specific range. For those simulations where each range was transformed by a unique value of k , we assumed gain normalization was applied within each range through its associated value of k . The resulting decision values (DV) were then computed as

$$DV = \frac{\sum_{i=1}^{10} \hat{X}}{g} \quad (4.2)$$

Finally, model choice probabilities were obtained by passing the obtained DV's through a sigmoidal response function (Equation 2.5, p. 46).

We then iterated over a grid of parameter combinations of late noise and

compression, limited as follows:

$$k \in \{e^{-2}, \dots, e^2\}$$

$$s \in \{0.001, \dots, 1.5\}$$

while keeping motor bias and leak constant at 0. For simulations with a single value of k , the parameter values were spaced over 101 steps, while the parameter values in simulations where k could vary freely for each range were spaced over 51 steps to reduce the computational cost.

4.3 Results

4.3.1 Model simulations

We started with the case where an agent employs a fixed compression of the number line x^k for any range. The agent can either treat numbers relative to their respective range (relative) or as lying on a single number line (absolute). We found that in both cases, under increasing levels of late noise, highest accuracy was achieved employing an anti-compression ($k > 1$; Figure 4.2). Moreover accuracy declined at a faster rate under absolute scaling and required greater levels of anti-compression to compensate for the increasing late noise (Figure 4.2B). Note that for relative scaling here, the simulations are in practice identical to the simulations Spitzer et al. (2017) reported in their work.

Alternatively, compression does not need to be fixed over ranges. It is possible an agent treats each range as a separate condition that affords its




Figure 4.2. Accuracy under late noise with fixed k . **(A)** Estimated model accuracy under relative scale. Orange dotted line represents maximum accuracy for any level of late noise. **(B)** Estimated model accuracy under absolute scaling. In both cases, anti-compression ($\log(k) > 0$) leads to higher accuracy under increasing levels of late noise. However, under absolute scaling performance drops more rapidly and demands greater compensation through anti-compression.

own compression. Independently, the number line can still be treated in relative or absolute terms. We therefore extended our simulations to allow for k to vary freely between the low and high range. Again anti-compression was optimal in all cases (Figure 4.3). When applying a relative scale, the simulations showed a linear relation between the optimal k of the low and high range (Figure 4.3B). This is equivalent to using a single k over all ranges. On the other hand, under absolute scaling the relation between the optimal k for low and high range was non-linear (Figure 4.3D). Anti-compression was greater for the high range compared to the low range, in particular for lower levels of late noise. Moreover, average accuracy declined much faster with increasing noise compared to relative scaling.




Figure 4.3. Accuracy under late noise with multiple k . **(A-B)** Estimated model accuracy under relative scaling. (A) shows for a subset of levels of late noise the heat maps of estimated accuracy. Orange points indicate maximum accuracy for each level of late noise. (B) displays estimated maximum accuracy for all simulated levels of s . Grey dots show projection of maximum accuracy onto the 2D parameter space of k . The optimal solution here is to maintain the same level of anti-compression over both ranges. **(C-D)** Same as above, simulating absolute scaling of the number line. (C) shows accuracy deteriorates much faster under increasing noise, while anti-compression remains the optimal solution. (D) The simulation predicts anti-compression should be stronger for the high range, especially under lower levels of noise. Cross markers show estimates that fell outside of the parameter space of k .

4.3.2 Interim discussion

These simulations lead to a number of predictions. First, if participants are trying to optimize performance given a certain level of late noise, they should display a deviation from linear processing of the number line towards anti-compression, as demonstrated by previous work (Luyckx et al., 2019; Spitzer et al., 2017). Additionally using a relative scale leads on average to higher levels of accuracy compared to an absolute scale. Our simulations further show that when using a relative scale, there is no benefit to employing multiple k 's per range. Thus, if participants perform 'optimally' in our task in the presence of late noise, they should adopt a relative scaling of the number line where the degree of anti-compression is the same for each range.

4.3.3 Model fits to human data

Participants were able to perform the task, with an overall average accuracy of 75% ($\pm$ SD 0.05). To determine whether (a) numbers were normalized within each range versus across the experiment and (b) the distortion of the number line differed between conditions, we performed a 5-fold cross-validation on 8 candidate models. In line with our predictions we found that models assuming absolute scaling showed overall worse cross-validation (mean LL = -207.15) compared to models assuming normalization within each range (mean LL = -198.82; Figure 4.4A). Within the latter class of models, a model with a fixed compression parameter k across all conditions best explained the data ($\text{pxp} = 0.66$; Figure 4.4B), as predicted by our simulations. Note that these results held even when only taking into account mixed blocks (mean $\text{LL}_{\text{abs}} = -205.64$, mean $\text{LL}_{\text{rel}} = -198.92$; $\text{pxp}_{\text{rel_k1}} = 0.99$), discarding the possibility that an absolute scale was used in those blocks where the range could vary from one trial to the next.




Figure 4.4. Model fitting results. **(A-B)** Using 5-fold cross-validation, we compared 8 candidate models under (1) relative (dark grey) or absolute (light grey) normalization and (2) with a varying number of k parameters. Both average negative log-likelihood estimates and protected exceedance probability favoured a model with relative scaling and a single fixed exponent k . **(C)** Probability density estimate of the distribution of k over all participants. The best fitting model showed a clear shift towards anti-compression ($\log(k) > 0$). Black vertical lines display reduced histogram. **(D)** Model predictions for each range from the best fitting model closely matched the empirical choice data. Bars and shaded area represent 95% confidence intervals.

Next, based on our simulations we predicted that anti-compression leads to better performance under late noise. We therefore set out to test if the winning model predicted values of k greater than 1 (= linear). Indeed we found that parameter estimates of k across participants showed a strong shift towards anti-compression of the number line ($t(20) = 5.56, P < 0.001$, two-tailed t-test) (Figure 4.4C). We then plotted qualitative fits of the best fitting model against the decision weights associated with each number within a range. Model estimates closely matched idiosyncratic patterns of human

choice data within each range, despite the fact that parameter estimates were fixed over the whole experiment for each participant (Figure 4.4D). If it is indeed the case that higher levels of late noise necessitate greater compensatory anti-compression, we would expect to find a negative relation between estimated late noise s and estimated compression k . In contrast to previous findings (Tsetsos et al., 2016; and see Chapter 5, Section 5.3.2) we could not find such a trade-off ($\rho = -0.22$, $P = 0.33$).

4.4 Discussion

Under the assumption that humans have limited cognitive resources to integrate information, we replicated the finding that participants in our sequential sampling paradigm non-linearly weighted the input. Model simulations showed that anti-compression of the number line in all cases led to improved performance in the task. Importantly, by using different ranges of numbers in different blocks, our paradigm allowed us to assess the contribution of normalization to robust processing. We found that normalization of the stimulus space leads to a slower drop in performance over different levels of late noise and requires less compensation through anti-compression. Finally, given a relative coding of the number space, only a single compression parameter was needed to achieve optimal performance. Based on a comparison of 8 different candidate models, we found that indeed participants normalized the number space and their choices were best explained by using a single compression parameter over all conditions. These conclusions held true even when only taking into account blocks where trials from the two ranges were mixed. It is perhaps surprising that participants did not use an absolute scale in mixed blocks and were able to flexibly switch between normalizing different ranges. First, by using distinct colour mappings for each range we

wanted to emphasise the different ranges and it is possible that participants were thus able to very quickly adjust to the coming number range. Second, there was never a direct comparison involved between numbers from the low and high range, therefore participants were never prompted to assess numbers as coming from a single number line.

Our findings further demonstrate that presenting two-digit numbers in the numerical integration task does not impair performance, nor lead to apparent differences in the asymmetric weighting of the number line. We chose the two adjacent ranges to maximize the opportunity for participants to represent the numbers on either an absolute or relative scale. If we had used two ranges that lay further apart, it might have been more likely for participants to ignore the possibility of an absolute scale.

Finally in our simulations we assumed that gain normalization was applied within each range. We acknowledge that more complicated forms of gain normalization could exist for mixed trials or for those instances with more than one compression parameter. Previous work has also shown that gain control can rapidly adapt, even at the single trial level (Cheadle et al., 2014). Given that model fits to the human data predicted a fixed compression and did not show different predictions when only taking into account mixed trials, we believe these assumptions to be sufficient for our simulations.

References

- Barlow, H. B. (1961). Possible principles underlying the transformations of sensory messages. In W. Rosenblith (Ed.), *Sensory communication* (pp. 217–234).
- Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not noisy, just wrong: The role of suboptimal inference in behavioral variability. *Neuron*, *74*(1), 30–39.
- Bhui, R., & Gershman, S. J. (2018). Decision by sampling implements efficient coding of psychoeconomic functions. *Psychological Review*, *125*(6), 985–1001.
- Carandini, M., & Heeger, D. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, *13*(1), 51–62.
- Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., DeGardelle, V., HerceCastañón, S., & Summerfield, C. (2014). Adaptive gain control during human perceptual choice. *Neuron*, *81*, 1429–1441.
- de Gardelle, V., & Summerfield, C. [C.]. (2011). Robust averaging during perceptual judgment. *Proceedings of the National Academy of Sciences*, *108*(32), 13341–13346.
- Drugowitsch, J., Wyart, V., Devauchelle, A. D., & Koechlin, E. (2016). Computational precision of mental inference as critical source of human choice suboptimality. *Neuron*, *92*(6), 1398–1411.
- Li, V., Castañón, S. H., Solomon, J. A., Vandormael, H., & Summerfield, C. (2017). Robust averaging protects decisions from noise in neural computations. *PLoS Computational Biology*, *13*(8), e1005723.
- Louie, K., & Glimcher, P. W. (2012). Efficient coding and the neural representation of value. *Annals of the New York Academy of Sciences*, *1251*(1), 13–32.

- Luyckx, F., Nili, H., Spitzer, B., & Summerfield, C. (2019). Neural structure mapping in human probabilistic reward learning. *eLife*, 366757.
- Michael, E., de Gardelle, V., Nevado-Holgado, A., & Summerfield, C. (2013). Unreliable Evidence: 2 Sources of Uncertainty During Perceptual Choice. *Cerebral Cortex*, 1–11.
- Moran, R., & Tsetsos, K. (2018). The standard Bayesian model is normatively invalid for biological brains. *Behavioral and Brain Sciences*, 41.
- Polania, R., Woodford, M., & Ruff, C. (2019). Efficient coding of subjective value. *Nature Neuroscience*, 22(January), 134–142.
- Sims, C. A. (2006). Rational inattention: Beyond the linear-quadratic case. *American Economic Review*, 96(2), 158–163.
- Spitzer, B., Waschke, L., & Summerfield, C. (2017). Selective overweighting of larger magnitudes during noisy numerical comparison. *Nature Human Behaviour*, 1, 1–8.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Group Bayesian Model Selection for Group Studies. *NeuroImage*, 46(4), 1004–1017.
- Summerfield, C. [Christopher], & Tsetsos, K. (2015). Do humans make good decisions? *Trends in Cognitive Sciences*, 19(1), 27–34.
- Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences*, 113(11), 3102–3107.
- van den Berg, R., & Ma, W. J. (2018). A resource-rational theory of set size effects in human visual working memory. *eLife*, 7, 1–31.
- Wei, X. X., & Stocker, A. A. (2015). A Bayesian observer model constrained by efficient coding can explain 'anti-Bayesian' percepts. *Nature Neuroscience*, 18(10), 1509–1517.

Wyart, V., & Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. *Current Opinion in Behavioral Sciences*, 11, 109–115.

Chapter 5

Selective integration during sequential sampling in posterior neural signals

5.1 Introduction

Biological brains evolved to be both precise and efficient. Responding accurately to external stimulation requires noisy sensory signals x to be transduced such that internal estimates $\hat{x}$ are as close as possible to their generative counterparts. Sequential sampling and integration allow the precision of sensory estimates to grow with the number of independent observations obtained (Wald & Wolfowitz, 1949), and neural signals in multiple brain regions implement a basic form of memory that allows optimal estimates to be gradually integrated during inference (Gold & Shadlen, 2007; Hanks & Summerfield, 2017). Forming more precise estimates of the sensory world is often consequential for reinforcement, both in the lab (e.g. when a monkey categorises a stream of noisy sensory signals in return for liquid reward) and

in the real world (e.g. when a consumer evaluates the quality of multiple relevant attributes of a product).

However, where information arrives in high volumes, the carrying capacity of the neural system can limit the precision of sensory estimates. Consider a binary discrimination judgment between two stimuli A and B , where information about their worth x_k^A and x_k^B arrives in parallel via a sequence of n discrete sample pairs k . Let us assume that decisions are made by integrating and comparing transduced sensory estimates under the corrupting influence of “late” noise, so that A is chosen if:

$$\sum_{k=1}^n [f(x_k^A - x_k^B)] + \mathcal{N}(0, \sigma) > 0$$

With limitless capacity ($\sigma = 0$) the best judgments will be made by simply comparing the summed estimates for A and B – in other words, the optimal $f(\cdot)$ is the identity function. Intuitively, if samples x are numbers and you are in possession of a calculator, the best thing to do is to simply compute their relative sum. However, under imperfect memory, comparative judgments of ground truth quality can paradoxically be promoted by more selective transduction that discards part of the sensory information (Li, Castañón, Solomon, Vandormael, & Summerfield, 2017; Moran & Tsetsos, 2018; Spitzer, Waschke, & Summerfield, 2017; Tsetsos et al., 2016). For example, when $\sigma > 0$, a accuracy-maximising strategy is to selectively integrate the locally highest valued sample (i.e. x_k^A when $x_k^A > x_k^B$) by partially down-weighting the less valued sample by a factor w ($1 > w > 0$). For illustration, imagine that deciding between two apartments to rent involves evaluating the candidates along multiple continuously-valued dimensions (price, location, size, etc.). As memory demands become prohibitive (e.g. as σ grows), a strategy that is partly based on tallying the number of dimensions

on which an alternative is higher valued (e.g. $w < 1$) will lead to more accurate choices. This is because ordinal (rank-based) strategies confer robustness on inference, just as non-parametric methods allow statistical calculations to be more robust to outlying data. Selective integration can be shown to be “optimal” because it optimises accurate responses given irreducible constraints on capacity. At the same time, however, selective integration is “irrational” in that it discards information and can lead to inconsistent choices, including violations of the normative axioms of transitivity and independence from irrelevant alternatives (Von Neumann & Morgenstern, 1944). Several models embodying this principle provide a good fit to human data in tasks involving sequential integration and comparison of discrete samples of information (Bhatia, 2013; Glickman, Tsetsos, & Usher, 2018; Gluth, Spektor, & Rieskamp, 2018; Summerfield & Tsetsos, 2015; Tsetsos, Chater, & Usher, 2012).

The framework of selective integration also makes predictions about neural signals, but these remain as yet untested. For example, several studies have shown that during sequential integration of decision information, the amplitude of single-trial EEG signals contralateral to the side of stimulation correlates with the strength of evidence conferred by each sample (Wyart, Myers, & Summerfield, 2015). Thus, we might expect that those samples that are selectively integrated by virtue of being the local “winner” would be encoded with higher gain, i.e. a steeper linear relationship between x and relevant EEG signals. Here, we tested this view by recording scalp electroencephalographic (EEG) data whilst participants performed a task that involved averaging and comparing bar height within two parallel streams.

5.2 Methods

5.2.1 Participants

Participants were 18 healthy adult volunteers with no history of neurological or psychiatric illness. Those who failed to complete both recording sessions ($N = 2$) or who performed at chance level ($N = 1$) were excluded. Analyses were conducted on the remaining 30 EEG sessions from 15 volunteers (N female = 8; N left-handed = 2; age = $24.29 \pm \text{SD } 4.48$). Sample size was not based on a power analysis, but rather preference was given to acquire a high number of data points per participant over two EEG recording sessions to achieve robust results (Smith & Little, 2018). The study was approved by the Oxford University Medical Sciences Division Ethics Committee (MSD/IDREC/C1/2009/1) and informed consent was given at the start of each recording session. Monetary compensation was based on performance during the experiment (approximately £25 per participant).

5.2.2 Experimental procedure

All stimuli were created in the Psychophysics 3 Toolbox (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007) for Matlab (MathWorks). The experiment was presented on a 17-inch CRT monitor with 60 Hz presentation rate at a viewing distance of ~ 70 cm. On each trial participants viewed a sequence of 9 pairs of bars of variable height (Figure 5.1A) that appeared in sequence, left and right of a central fixation dot. The task framing differed across sessions. In one recording session participants were asked to indicate with a button press the stream with the highest average bar height ('high frame') and in the other session they were asked to indicate the stream with the lowest average height ('low frame'). The order of the two sessions was counterbalanced

over participants and both sessions were separated by at least one week. We excluded participants who did not complete both sessions, because the framing manipulation allowed us to cleanly orthogonalize perceptual value (the raw bar height in pixels) from decision value (the level of evidence in favour of either response). The instructions stated that the bars indicated the value of two stock options whose prices fluctuated over time. They were asked to buy the most favourable (high frame) or sell the least favourable (low frame) stock option in different sessions.

Each trial started with a white fixation dot (radius = 5 pixels [px]) presented centrally on a grey background. The dot remained on screen for the whole duration of the trial. All bar stimuli were presented within two black rectangular placeholders (60 by 200 px; visual angle $\sim 4^\circ$) on either side of the fixation dot at a horizontal distance of 160 px. Each pair of bars remained on screen for 350 ms, and between successive pairs of bars, there was a jittered gap of empty black frames, lasting on average 150 ms (uniformly drawn between 100 – 200 ms). The first bar was a forward mask (300ms after fixation onset) of maximal bar height that occurred on both sides and was not included in the analysis. Bars were thus presented at a rate of ~ 2 Hz (including the jitter) to minimise steady state neural responses. When all 9 bar pairs had been presented, a backward mask appeared again on screen for 50 ms. The fixation dot then turned black, indicating to the participants that they could respond by pressing the A and L keys on a QWERTY standard keyboard using the left and right hand to choose left and right stream respectively. Response mapping was fixed for the entire experiment. If participants failed to respond within 3 seconds, the fixation dot would turn red for 1000 ms and the words ‘Too late’ appeared in red above the fixation dot. When they responded before the deadline, the fixation

dot would change colour for 500 ms: green for a correct response and red for incorrect. The next trial started after an inter-trial interval (ITI) of 300 ms. A trial could thus last for a maximum of 9350 ms (assuming 3 seconds with no response).

Trials were drawn from one of four bar height distributions (conditions), randomly intermixed throughout the experiment. In two conditions ('Low variance' and 'High variance'), samples for the left and right stream were drawn pseudorandomly from a Gaussian distribution whose mean μ was in turn drawn from a uniform distribution between 80 px and 120 px, and whose standard deviation σ set to 10 (low variance) or 20 (high variance conditions). We used a resampling method to ensure that sample mean and variance closely matched the generative means and variance. To produce a correct answer, 6 px were subtracted from one of the two stream means.

We also added two more conditions which permitted greater control over the number of winning samples (i.e. those samples with locally greater decision evidence) on either side of the stream, allowing us to manipulate the correct answer independently from the number of winning samples contained within that stream. Samples for these two conditions, 'Frequent winner' and 'Infrequent winner', were generated as follows. In the frequent winner condition, the correct stream was manipulated to contain the winning sample in 2/3 of the sample pairs, while in the infrequent condition the correct stream only contained the winning sample 1/3 of the time. To achieve this, the mean bar height (M) of each trial was first drawn from a uniform distribution ranging between 110 and 130 px. Next, values were calculated separately for every three sample pairs ('triplets' T) in the trial. A deviation (δ) from the trial mean for each triplet pair T was drawn from a uniform distribution $\delta \sim U(15, 25)$. Additional noise ϵ_1 and ϵ_2 was drawn per triplet T from a




Figure 5.1. Task design and model schematic. **(A)** Participants viewed a sequence of nine pairs of bars varying in height and compared the average bar height of the left and right stream. In two separate recording sessions, participants indicated the sequence with either the highest or lowest average bar height. Numbers in bold below frames indicate duration of each frame in milliseconds. **(B)** The selective integration model assumes a separate accumulation process for each option, here the left and right bar sequence. Two concurrent input values X^A and X^B are transformed with a factor w based on their relative difference, down-weighting the lower input value and up-weighting the higher input value. All transduced values I^A and I^B are then integrated over time, with greater loss of information for earlier samples. Finally, the model generates a choice probability based on the difference in integrated values between the two streams.

normal distribution $\epsilon \sim \mathcal{N}(0, 3)$. For every triplet T , the two streams T_A and T_B would have the following form:

$$T_A = \{M + \epsilon_1, M + \delta + \epsilon_2, M + 2\delta + (\epsilon_2 - \epsilon_1)\}$$

$$T_B = \{M + \delta, M + 2\delta, M\}$$

In this form, T_B always contains two winning samples ($M + \delta$ and $M + 2\delta$), while the means of both triplets are identical. To produce a correct answer, 6 px were subtracted from one of the two stream means. In the frequent winner condition the mean of stream A, which contained most of the losing samples, was adjusted to make it the incorrect answer, while in the infrequent winner condition the mean of stream B, containing the most winning samples, was adjusted. Adjustments in low and high frames were inversions of one another.

To prevent participants from shifting gaze directly at one of the two streams, an area around the fixation dot (256 x 205 px; visual angle $\sim 2^\circ$) was defined in which participants had to keep fixation. Eye movements were monitored online using a Tobii EyeX eye-tracker (Tobii Technology, Stockholm, Sweden) and when the eyes moved outside the fixation area, the trial was classified as incorrect and the message ‘Eyes moved!’ appeared above the fixation dot. These trials were also omitted from analysis (0.35%).

Each recording session consisted of 600 trials, divided in 10 blocks, resulting in a maximum of 1200 trials (10 800 samples) for each participant. Participants could take self-timed breaks in between blocks. One session lasted around 2.5 h, including preparation of the cap, 60 minutes of task performance and removal of the cap at the end of the experiment.

5.2.3 Selective Integration model

The model decision is based on the output of two accumulators Y^A and Y^B that integrate the input values of the left and right stream respectively (Figure 5.1B). Input values X^A and X^B are the raw pixel heights H of each sample (bar), inverted for the low frame ($200 - H$). Each integrator is updated separately for each sample k according to the following formula:

$$\begin{aligned} Y_k^A &= Y_{k-1}^A + I_k^A \cdot \lambda \\ Y_k^B &= Y_{k-1}^B + I_k^B \cdot \lambda \end{aligned} \quad (5.1)$$

Y^A and Y^B are both initialized to zero. I^A and I^B are the transformed values of the input after graded selective integration:

$$\begin{aligned} I_k^A &= (1 - w) \cdot X_k^A \\ I_k^B &= w \cdot X_k^B \end{aligned} \quad \text{with} \quad \begin{cases} w < 0.5 \text{ where } X^A > X^B \text{ for sample } k \\ w > 0.5 \text{ where } X^A < X^B \text{ for sample } k \end{cases} \quad (5.2)$$

Where w is a gating variable that is determined by passing the difference between input values ($\Delta X = X^B - X^A$) through a logistic function

$$w = \frac{1}{1 + e^{-m(\Delta X)}} \quad (5.3)$$

with the slope of m determining the extent to which the difference between the two input values attenuates the input values X . Values of m close to zero indicate a limited sensitivity to the difference in bar height, treating the gating of locally ‘winning’ and ‘losing’ samples almost equally when the difference between them is small versus when their difference is large (note that the ‘winning’ sample changes with the frame of reference i.e. whichever

sample is highest in the high frame and lowest in the low frame). Larger values of m then indicate an increased sensitivity to the magnitude of the difference between the bar heights in a sample pair, giving rise to a tendency to more strongly overweight the locally winning sample for larger differences; negative values of m have the converse effect. The transfer function itself ensures that for very large differences between the bars, where it is clear which of the two bars wins, the winning input value is integrated close to its original value, while the losing sample value is almost entirely suppressed (i.e. w is close to 0 or 1). On the other hand, for input values that lie close together and comparison is difficult, both values carry approximately equal weight to their accumulators. A recent study demonstrated that a ‘graded’ selective integration such as the one proposed here outperformed a ‘binary’ selective integration that works independent of the size of the sample difference, both in predicting human choices and its robustness to noise (Glickman et al., 2018).

The model further assumes a leaky accumulation, with earlier samples carrying less weight on the final decision. Each transduced sample is therefore multiplied with the leak function value:

$$\lambda = \frac{(1-l)^{n-k}}{(1-l)} \quad (5.4)$$

with l the amount of information loss and n the total number of samples.

Finally, the model outputs a choice probability based on the difference between the accumulators after all evidence has been accumulated:

$$P(B) = \kappa + \frac{1-2\kappa}{1+e^{-s(Y^B-Y^A)}} \quad (5.5)$$

s is the slope of the response function, often referred to as ‘late’ or

‘integration noise’, a factor that adds uncertainty to the final choice. Finally, a lapse rate parameter κ was added to the response function to allow for higher error rates in the model.

5.2.4 Model fitting procedure

The model has four free parameters: selective gating parameter m , leak l , late noise s and lapse rate κ . Best fitting parameter sets for each participant were obtained by minimizing the negative sum of the log-likelihoods using a scatter-search based global optimization solver in Matlab. The parameter search space was constrained as follows:

$$m \in \{-0.2, \dots, 0.2\}$$

$$\kappa \in \{0.001, \dots, 0.2\}$$

$$l, s \in \{0.01, \dots, 0.5\}$$

To test whether the gating parameter m significantly contributed to the model performance, two models were compared: the full SI model and a fixed gating model. For the fixed gating model m was set to 0, which effectively gave equal weight to each input ($w = 0.5$), whereas for the full model m was unrestricted, allowing w to take any value between 0 and 1. A second model comparison tested the full model against a simple counting strategy of winners, to assure that the parametric values of the samples were taken into account during evidence integration. For the counting model, parametric values were ignored and replaced by 1 and 0 for winning and losing samples respectively. Model fits were compared through split-half cross-validation, a method that allows comparing models with different numbers of free parameters through

their generalizability to new data. One recording session would serve as the training set and the best fitting parameters obtained from the training set were then used to estimate how well they predicted responses in the other session and vice versa. Negative log-likelihoods were summed per participant and fed into a Bayesian Model Selection procedure (Stephan, Penny, Daunizeau, Moran, & Friston, 2009) to compute posterior evidence for one model over another. Best fitting parameters for the model used in subsequent analyses were estimated from the collapsed data over both recording sessions.

5.2.5 Behavioural analyses: regression models

All analyses of human data were performed on the collapsed data from both recording sessions, with input values for the ‘low frame’ session inverted ($200 - H$) to reflect the momentary decision value X rather than raw bar height H . The impact of winning and losing samples (i.e. the bars with respectively the highest and lowest decision evidence in a sample pair) on choice was tested through a generalized linear model (GLM), modelled with a binomial distribution and a logit link function, predicting the probability of choosing the right stream, as follows:

$$P(R) = \beta_0 + \beta_1 \cdot \sum_{k=1}^9 X_k^L(L > R) + \beta_2 \cdot \sum_{k=1}^9 X_k^L(L < R) + \beta_3 \cdot \sum_{k=1}^9 X_k^R(L < R) + \beta_4 \cdot \sum_{k=1}^9 X_k^R(L > R) + \epsilon \quad (5.6)$$

The four parametric regressors in the model coded for the cumulative sum of input values for all (1) winning samples on the left side [$X^L(L > R)$], (2) losing samples on the left side [$X^L(L < R)$], (3) winning samples on the right

side $[X^R(L < R)]$ and (4) losing samples on the right side $[X^R(L > R)]$. The regressors were always standardized before parameter estimation. Parameter estimates for the left stream were then sign-flipped and averaged with the estimates of the right stream, in order to obtain an average modulation of winning and losing samples.

The SI model further allows for a ‘recency effect’ through the leak parameter, whereby earlier samples carry less weight in the final choice than later samples. We tested this assumption through a new GLM predicting choices for the right stream based on the signed difference of each sample pair (right minus left).

$$P(R) = \beta_0 + \beta_1 \cdot \Delta(X_1^R, X_1^L) + \dots + \beta_9 \cdot \Delta(X_9^R, X_9^L) + \epsilon \quad (5.7)$$

The model thus consisted of 9 parametric regressors representing the serial position in the stream. If choices were driven more by recent samples, coefficient estimates should be higher for regressors coding for those later samples.

All statistical tests were performed at the group level. Given the relatively low number of subjects, we opted for non-parametric tests that do not assume a normal distribution of the data.

5.2.6 EEG acquisition

EEG signals were recorded using a Neuroscan system with SynAmps-2 digital amplifiers with 60 Ag/AgCl electrodes located at FP1, FPz, FP2, F7, F5, F3, F1, Fz, F2, F4, F6, F8, FT7, FC5, FC3, FC1, FCz, FC2, FC4, FC6, FT8, T7, C5, C3, Cz, C2, C4, C6, T8, TP7, CP5, CP3, CP1, CPz, CP2, CP4, CP6, TP8, P7, P5, P3, P1, Pz, P2, P4, P6, P8, PO7, PO5, PO3, POz, PO4, PO6, PO8, O1, Oz, and O2. Four additional EOG electrodes in dipolar montage

(two horizontal, two vertical) were recorded, together with one mastoid for reference. Electrode impedance was brought below 10 k Ω before recording. Data was collected at a sampling rate of 1 kHz and high-pass filtered at 0.01 Hz.

5.2.7 EEG pre-processing

Pre-processing was done in Matlab using functions from the EEGLAB toolbox (Delorme & Makeig, 2004) and custom scripts. First, data was downsampled to 250 Hz, subsequently low-pass filtered at 40 Hz and then high-pass filtered at 0.5 Hz. Excessively noisy channels were identified through visual inspection for each participant and interpolated based on the weighted average of the surrounding electrodes. Next the data was re-referenced offline to average reference (excluding EOG channels). Trial epochs were extracted spanning 1 second prior to the fixation dot onset until 7 seconds after. Epochs were subsequently baselined relative to the pre-fixation time window of -500 to 0 ms. Artefacts related to eye-blinks and other sources of consistent noise were identified through Independent Component Analysis (ICA) and removed from the data after visual inspection. For the sample-based regressions, the data were epoched relative to each sample pair onset, starting at 100 ms before until 750 ms after sample onset and baselined again relative to the full pre-stimulus window, to exclude any systematic offset during the length of the trial.

HEOG channel data were pre-processed using a similar pipeline, apart from average referencing and ICA.

5.2.8 Time-frequency transformation

The pre-processed epochs spanning 8 seconds were transformed into the time-frequency domain using functions from the Fieldtrip toolbox (Oostenveld, Fries, Maris, & Schoffelen, 2011) for Matlab and custom scripts. Power was calculated every 25 ms within an epoch by convolving a sliding Morlet wavelet (width = 7 cycles), for frequencies between 8 and 38 Hz, incremented in steps of 3 Hz. To allow for a comparison of power estimates over frequency bands, each frequency within an epoch was dB transformed relative to the baseline window [-300 300] ms, a ‘silent’ period including the ITI and start of a new trial, using the following formula:

$$dB = \log_{10}\left(\frac{signal}{baseline}\right) \quad (5.8)$$

The data was subsequently epoched into smaller epochs between [-100, 750] ms relative to sample pair onset. For visualization purposes the time-frequency results were interpolated, while cluster statistics were performed on the raw data.

5.2.9 EEG analyses: regression models

To understand how EEG signals were modulated by the stimulus information, we used a regression-based approach. Single-trial EEG data was regressed against parametric predictors for each time-point and electrode independently. This analysis tests for a linear relationship between the magnitude of decision values and the momentary amplitude of the EEG signal at each time point following each sample. Coefficient estimates obtained from the regression model reflect the slope of the modulation, i.e. the strength of the linear relationship. Group-level statistics were subsequently computed over the

obtained time courses of parameter estimates over all participants.

First, we constructed a sample-based linear regression model that allowed us to dissociate the neural modulation of winning and losing samples:

$$\begin{aligned}
 EEG_k = & \beta_0(L > R) + \beta_1(R > L) + \beta_2 \cdot |X_k^L| \mid (L > R) + \beta_3 \cdot |X_k^L| \mid (L < R) + \\
 & \beta_4 \cdot |X_k^R| \mid (R > L) + \beta_5 \cdot |X_k^R| \mid (R < L) + \\
 & \beta_6 \cdot |HEOG| \mid (L > R) + \beta_7 \cdot |HEOG| \mid (L < R) + \epsilon
 \end{aligned}
 \tag{5.9}$$

with β_0 and β_1 as binary indicator variables for samples where the left (L) or right (R) bar in the pair had the highest decision value respectively. The four parametric regressors $\beta_2 - \beta_5$ code for the absolute decision evidence (X_k) of the bar when the left bar won (β_2), left lost (β_3), right won (β_4) and right lost (β_5). Effectively, each sample pair related to four regressors (β_0 , β_1 and two of the four parametric regressors). The trial mean was subtracted from all bars to eliminate any differences related to the absolute size of the bars. The HEOG signal was added as a nuisance regressor to account for variability related to small eye movements. The single trial values of the HEOG channels were calculated based on an averaged time window of 100 ms centred around the negative peak of the parametric regressors (176 - 276 ms) from an initial model that did not contain the HEOG nuisance regressors. Single-trial EEG signal locked to stimulus onset was subsequently regressed against the full model with 8 regressors, independently for each sample pair. Estimated time courses were averaged over all sample pairs before statistical inference.

We defined two regions of interest (ROI) based on previous research (Wyart et al., 2015) to test for early modulation of sample evidence: left

(O1, PO5, PO7, P5, P7, CP5, TP7) and right (O2, PO6, PO8, P6, P8, CP6, TP8) occipito-parietal (OP) electrodes. The coefficients β_0 and β_1 in our regression model encode the average EEG signal for sample pairs where the left (β_0) or right sample won (β_1). To compare categorical modulation of stimulus-evoked activity, we averaged the activity over those electrodes contralateral to the winning sample, namely right OP electrodes for β_0 (trials where the left sample won) and left OP electrodes for β_1 (trials where the right sample won). This is equivalent to flipping the scalp maps and treating each sample pair as if the right sample won. The opposite hemispheres then contained the neural signals modulated by the losing sample. Time series for a parametric modulation of winning samples (X_{win}) were obtained in a similar way: averaging over the right OP electrodes when left won (β_2) and left OP electrodes when right won (β_4). Time series for parametric modulation of losing samples (X_{lose}) were obtained at right OP electrodes when left lost (β_3) and left OP electrodes when right lost (β_5). Significance of adjacent time points in the averaged time series were tested using non-parametric cluster-based permutation tests (cluster-defining threshold and corrected significance level at $P < 0.05$, 5000 iterations) (Maris & Oostenveld, 2007). Estimated time series were smoothed per subject using a low-pass filter (Butterworth filter, cut-off = 25 Hz), removing small fluctuations in the signal that could obscure larger clusters from the permutation test.

The same regression model and contrasts were used to study differences in lateralization of alpha band activity, with an additional iteration over frequency bands. The averaged time window was temporally smoothed with a 100 ms full-width at half-maximum (FWHM) Gaussian kernel before cluster-based permutation test (cluster-defining threshold and corrected significance level at $P < 0.025$). We used a lower p-value threshold for our cluster

correction algorithm for the time-frequency analysis, to avoid clusters from different frequency bands to merge into one, potentially inflating the reliability of smaller clusters, a well-known problem with cluster-based analysis methods (Maris & Oostenveld, 2007).

Next, we sought to identify neural signals related to evidence accumulation. We formalized ‘accumulated evidence’ as the difference between the cumulative sum of evidence of the right minus left stream. The model tested evidence integration during the course of the trial and was therefore conducted at the sample level (for samples 2 to 9), regressing EEG signals from sample k against the absolute sum of differences up to the previous sample (i.e. 1 to $k - 1$) and controlling for the evidence on sample k :

$$EEG_k = \beta_0 + \beta_1 \cdot \sum_{i=1}^{k-1} |\Delta(X_i^R, X_i^L)| + \beta_2 \cdot |\Delta(X_k^R, X_k^L)| + \epsilon \quad (5.10)$$

Since the outcome is rectified, i.e. unsigned with respect to the response hand, this measure was not confounded with motor responses. Because the cumulative difference for early samples was too highly correlated with the generative data, making it hard to disentangle the two, we opted to derive the independent variables from the model estimates rather than the generative data. The cumulative values included in the regression model therefore represented a leaky integration of weighted samples rather than a simple addition of raw input values.




Figure 5.2. Behavioural results and model fits. **(A)** Parameter estimates from a regression model predicting the probability of choosing the right-hand sequence, both in humans (blue/green) and the fitted SI model (red). In line with predictions of the SI model, sample values were weighted higher on average when decision evidence was greater relative to the concurrent sample. A SI model fitted to the choice data was able to recreate this pattern. **(B)** SI predicts that a tendency of tallying winning samples will lead to lower accuracy in trials where the incorrect sequence contains most winners. Performance was indeed lower in the ‘infrequent winner’ condition, but not significantly so from the ‘frequent winner’ condition. Qualitatively, the full SI model (red) resembled human accuracy better in all four conditions compared to a model where the gating parameter was fixed to 0. Black vertical lines show individual participant’s accuracy and grey shaded area represents non-parametric kernel estimate (bandwidth = 0.03) of the accuracy distribution in our sample. [continued on the next page]

(C) Parameter estimates from a new regression model showed a ‘recency effect’, where earlier samples carried less weight on participants’ final choice. The fitted SI model closely followed the same trend (red). Error bars and shaded area represent 95% confidence interval. (D) The slope of the response function (s) in the fitted SI model correlated significantly with the slope of the selective gating function (m), suggesting participants with greater levels of late noise compensate through stronger gating of information. Grey line represents least-squares fit line.

5.3 Results

5.3.1 Behavioural results

We first examined the behavioural data to test whether participants weighted the samples as described by the selective integration model. We fitted a logistic regression model that separately estimated the influence of winning and losing samples on right-hand choices (Equation 5.6). The null hypothesis is that there will be no difference in how much influence winning samples (e.g. X^A when $X^A > X^B$) and losing samples (e.g. X^B when $X^A > X^B$) carry on choice. However, we found that parameter estimates for winning samples (Mdn $\beta_{win} = 3.12$) were significantly higher than for losing samples (Mdn $\beta_{lose} = 3.06$; $Z = 3.41$, $P < 0.001$, two-sided Wilcoxon paired signed rank test), suggesting there is indeed a difference in how samples were weighted according to their relative decision value (Figure 5.2A).

If participants give more weight to winning samples, their performance might be poorer in a condition where the incorrect stream contains the most winning samples, a so-called ‘frequent-winner’ effect (Tsetsos et al., 2016). To maximize the opportunity to compare trials where the winning stream contained the most versus the fewest winning samples, we included two

conditions that explicitly contained these types of trials (‘frequent winner’ and ‘infrequent winner’) as well as two conditions where the height variance was either low or high. A non-parametric omnibus test (Friedman test) first confirmed there were differences in choice accuracy between the four conditions ($X^2 = 13.34$, $P = 0.004$; $d.f. = 3$, $N = 15$) (Figure 5.2B). We then ran a post hoc pairwise comparison between all conditions (two-sided Wilcoxon signed rank test) and corrected for multiple comparison using FDR correction (Benjamini & Hochberg, 2009). Unlike previous reports (Tsetsos et al., 2016), we did not observe a statistically significant difference between the ‘infrequent winner’ (Mdn = 0.66) and ‘frequent winner’ conditions (Mdn = 0.76; $Z = 1.62$, $P = 0.1054$), although the results did show a numerical trend in the predicted direction and a selective integration model fitted to the human choice data predicted a difference between the (in)frequent winner conditions (see below). Participants performed significantly worse in the ‘infrequent winner’ condition, but only in comparison to the ‘low variance’ condition (Mdn = 0.77; $Z = 2.24$, $P = 0.0248$). Participants also performed significantly worse in the more difficult high variance condition (Mdn = 0.72) compared to the low variance condition ($Z = 3.14$, $P < 0.001$).

Finally, the SI model predicts a ‘recency effect’, where samples presented later in a trial should carry more weight on the final decision, because there was less time for this information to be lost. A new logistic regression model predicting right-hand choices (Equation 5.7) showed that indeed, when sample pairs were assessed based on their serial position in a trial, parameter estimates increased over time, indicated by a significantly increasing slope fitted to each subject’s parameter estimates (Figure 5.2; Mdn slope = 0.13; $Z = 3.41$, $P < 0.001$).

5.3.2 Model fits to human data

A more formal test for selective integration can be achieved by fitting the SI model to human data and comparing it to an equivalent alternative model that does not show selective integration. A full SI model with 4 free parameters (gating m , leak l , late noise s and lapse rate κ) was fitted to each participant's trial-wise choice data ($m = 0.0377 \pm SD 0.03$; $l = 0.234 \pm SD 0.1$; $s = 0.0414 \pm SD 0.03$; $\kappa = 0.0176 \pm SD 0.02$; Table D.1). To assess whether the gating parameter meaningfully contributed to the model fits, we compared the full model to a fixed gating model where the gating parameter m was fixed to 0, eliminating the unequal weighting of sample inputs. Both models were tested through two-fold cross-validation, estimating parameter fits on one recording session and testing model predictions on the other session. Negative log-likelihood (LL) estimates of the test sets were then summed per participant and aggregate LL were fed into a Bayesian Model Selection procedure (Stephan et al., 2009) to determine the best fitting model. The full model (mean LL = -623.57) was favoured over the fixed model (mean LL = -627.77) with an exceedance probability of 65% suggesting the gating parameter of the SI model meaningfully contributed to explain choice behaviour. To determine qualitative fits of both models, we reran our behavioural analyses on model choices. The full model was able to reproduce the difference in parameter estimates for winning and losing samples (Figure 5.2A) and the recency effect (Figure 5.2C), even though the model was not specifically fit to these data points. The estimated weighting parameter m from the full model also significantly correlated with the differences in regression weights across participants, suggesting the observed difference in regression weights is likely linked to selective weighting (Spearman $\rho = 0.64$, $P = 0.012$). Only the full model was further

able to capture the lower performance in the ‘infrequent winner’ condition (Figure 5.2B), and in general captured the patterns in human performance per condition qualitatively better than the fixed model.

To ensure that a simple counting strategy of winners was not able to predict human choices as well as the full SI model, we repeated the same cross-validated model comparison as above with a counting model and the full model. The counting model performed much worse than the full model indicated by the much lower log-likelihood estimates (mean LL = -766.90) and an exceedance probability of the full model of 100% over the counting model.

Finally, Tsetsos et al. (2016) showed that the fitted parameters for gating and late noise were highly correlated in their data set (but not for simulated data), suggesting that the strength of the gating process compensates for higher levels of integration noise. We replicated this finding in our data set (Spearman $\rho = -0.91$, $P < 0.0001$) (Figure 5.2D). To ensure that this correlation was not due to a trade-off between the two parameters (m and s), we performed parameter recovery on simulated parameter combinations while systematically varying either m or s and assessed how well our fitting procedure could recover the original parameters (Appendix D). We could successfully recover the original parameters, while the recovery error was not significant for any of the remaining parameters (all $P > 0.115$, Figure D.1). Moreover, we could not find a correlation between m and s in either recovery procedure (all $r < 0.0001$, all $P > 0.99$).

5.3.3 Early modulation of posterior EEG activity

The SI model states that selective gating occurs at the level of the individual samples before evidence is passed to a subsequent integration stage. We thus




Figure 5.3. Early modulation by relative sample evidence in posterior electrodes. A linear regression model was used to separately assess the influence of winning and losing samples on EEG signals at each time point, electrode and sample pair. Coefficients for left sample regressors in our model were horizontally flipped, so winning samples were projected to the left hemisphere and losing samples to the right hemisphere. Plots represent coefficient estimates averaged for all nine sample pairs. **(A)** Coefficient estimates for the parametric regressors in our model showed that winning and losing samples were both significantly encoded in the EEG signal around 250 ms (bottom coloured lines, $P_{cluster} < 0.05$) in posterior electrodes (coloured dots on the scalp plots) contralateral to the sample location. More importantly, winning samples were encoded significantly stronger compared to losing samples (grey shaded area, $P_{cluster} < 0.05$). Shaded coloured area indicates SEM. Top and bottom scalp plot show significance of coefficient estimates for losing and winning samples respectively at the identified time window of significant dispersion. **(B)** By contrasting the intercept coefficients of the regression model, we observed a small additive effect that was stronger for winning samples compared to losing samples, but only when averaging between 200 and 300 ms after stimulus onset in posterior electrodes. The orange line represents the difference between the time series of the two intercepts. Scalp plot shows the left minus right hemispheric difference of coefficient estimates. Shaded coloured area represents SEM.

(C) An identical regression model was used to explain time-frequency data. We did not find a significant difference in coefficient estimates between winning and losing samples based on their parametric values. Colour map shows difference between winning and losing samples at the electrodes indicated on the scalp plot. Scalp plot shows difference between winning and losing samples and hemispheres at same time window as (D) for the 11 Hz frequency band. Time-frequency data was interpolated for visualization purposes, while cluster statistics were performed on the raw data. (D) However, we did observe an additive effect in alpha frequencies (8 – 12 Hz), with greater suppression for winning samples compared to losing samples in the same contralateral posterior electrodes starting around 375 ms after sample onset ($P_{cluster} < 0.025$). Colour map shows the difference between winning and losing samples at the electrodes indicated on the scalp plot. Scalp plot shows difference between winning and losing samples and hemispheres at the time of the identified cluster for the 11 Hz frequency band.

predicted that posterior EEG signals would encode locally “winning” samples with higher gain than “losing” samples, i.e. that the slope of regression linking decision information to EEG amplitudes would be steeper for winners than losers. In addition to this multiplicative effect, we also tested for an additive bias, i.e. that EEG signals were higher contralateral to the winning sample, irrespective of its input value.

Since samples were presented parafoveally, we expected these effects to occur contralateral to the location of each sample. We therefore focused our analyses on two a priori defined posterior regions of interest (ROI): left and right occipito-parietal electrodes. We constructed a linear regression model (Equation 5.9) with two intercepts, coding whether the left or right sample won within a sample pair (to test the additive effect), and four parametric regressors coding the sample evidence separately for left and right and whether the sample won or not (to test the multiplicative effect).

We first examined the coefficients associated with the parametric regressors that, at each time point and electrode, reflect the slope of the relationship between the magnitude of the decision evidence and EEG amplitude, separated for winning and losing samples. After collapsing over posterior electrodes contralateral to winning or losing samples, we found that the sample evidence of both winning and losing samples was encoded in the EEG signal around 250 ms, but more importantly, this was significantly stronger for winning samples ($P_{cluster} < 0.05$; Figure 5.3A). This suggests that, as predicted, samples carrying equal decision evidence are encoded more strongly in the EEG signal when they are the winning sample in a sample pair. Furthermore, when using the model transformed sample values as inputs to the regression, the difference between the two time series was no longer significant, indicating that the observed difference in EEG signal modulation was in fact related to the differential weighting in the model (Figure D.3).

The intercepts of the regression model then reflect the average signal in trials where either the left or the right sample won, independent of their decision evidence. At each time point and electrode, the parameter estimates reflect an additive effect of sample identity (winner or loser) on EEG signals. To test a difference in the additive effect of winning and losing samples, we averaged intercept parameter estimates at ROI electrodes contralateral to all winning samples and compared them to the averaged intercept parameter estimates at ROI electrodes contralateral to all losing samples. This initially yielded a significant effect whereby signal in posterior electrodes responded to winning samples more strongly than losing samples ($P_{cluster} < 0.05$; Figure 5.3B) between ~ 200 and 300 ms after sample onset. However, when adding HEOG nuisance regressors, the effect was attenuated and no longer

survived cluster correction ($P_{cluster} > 0.05$). It did, however, remain robustly significant when averaging over a 200 – 300 ms time window ($Z = -3.85$, $P = 0.0001$, Wilcoxon signed rank test). We thus interpret the additive effect with caution; there is a possibility that small eye movements could at least partially explain the neural effects we find here.

The negative contralateral modulations reported here show a strong resemblance to the well-known N2pc component from event-related potential (ERP) literature (Luck, 2012; Luck & Hillyard, 1994), a posterior negativity contralateral to the (covertly) attended target starting around 200 ms after stimulus onset. For the additive effect, this is expected, because the analysis we conducted on the intercept is closely related to the computation of an ERP. However, the fact that this signal is also modulated by the decision information nuances and extends previous work with trial averages.

5.3.4 Alpha suppression

It has previously been suggested that SI could occur when selective attention is oriented to the winning sample (Glickman et al., 2018). Previous work has reported that a lateralized suppression in alpha (8 – 12Hz) power occurs when participants orient visuospatial attention towards a contralateral target (Bacigalupo & Luck, 2019; Sauseng et al., 2005). We used the same sample-based regression model as above (Equation 5.9), but now predicting sample-wise fluctuations in power over frequency bands between 8 and 38 Hz. The same occipito-parietal contrasts for the intercepts of the model revealed a significant cluster of greater alpha suppression contralateral to the winning sample compared to the losing sample from ~375 ms onwards ($P_{cluster} < 0.025$) (Figure 5.3D). Note that these effects survived the inclusion of horizontal eye movements and might imply that covert attention was




Figure 5.4. Encoding of evidence integration in central electrodes. We found a significant positive relationship with the accumulated evidence of $1 : k - 1$ in central electrode Cz around 300 ms after stimulus onset. Scalp plot shows EEG activity averaged over the significant cluster around 300 ms. Bottom coloured lines indicate significant clusters ($P_{cluster} < 0.05$). Coloured shaded area indicates SEM.

dynamically shifting to the winning sample. We could not find any statistical evidence for a parametric modulation of alpha suppression by input value ($P_{cluster} > 0.05$) (Figure 5.3C).

5.3.5 Evidence accumulation

In the SI model, samples are first weighted according to their relative value, and then integrated into a cumulative signal that eventually determines choice. Previous studies have observed in both perceptual and value-based decision making tasks that central EEG signals build up in concert with the strength of accumulated evidence (Donner, Siegel, Fries, & Engel, 2009; O’Connell, Dockree, & Kelly, 2012; PISAURO, Fouragnan, Retzler, & Philiastides, 2017; Twomey, Murphy, Kelly, & Connell, 2015; van Vugt, Simen, Nystrom, Holmes, & Cohen, 2012; von Lautz, Herding, & Blankenburg, 2019). We examined

encoding of accumulated evidence over the course of the trial with a regression model that tested whether the cumulative difference in absolute (unsigned) decision information up to sample $k - 1$ was encoded in neural signals following the onset of sample k (Equation 5.10). Accumulated evidence (formalized as the leaky integrated values from the fitted SI model) was significantly encoded around 250 – 400 ms in the proximity of central electrode Cz ($P_{cluster} < 0.05$; Figure 5.4) suggesting neural signals were tracking the strength of evidence up to the current sample. Although the focus of these signals is slightly more frontal than previously reported, they resemble the “centro-parietal positivity” or CPP reported by O’Connell et al. (2012), which has also been related to the updating of the decision evidence.

5.4 Discussion

In this report we tested for, and obtained, neural evidence in support of a “selective integration” policy during human decision-making. Input values conferred by each of two samples (bars) was weighted differently according to its local rank, and this behavioural effect was accompanied by modulation of neural signals over posterior electrodes. This modulation arose early after sample onset (~ 250 ms), while signals encoding the strength of accumulated evidence peaked later in time (~ 300 ms) and more centrally.

Selective attention has been hypothesized to be an important driving mechanism, for selective integration models (Glickman et al., 2018) and related theoretical accounts (Bhatia, 2013; Gluth et al., 2018). Our results seem to suggest an important contribution of selective attention in multiple ways. First, contralateral alpha power suppression has long been related to visuospatial attention towards a target (Bacigalupo & Luck, 2019; Sauseng et al., 2005). We found that 375 ms after each stimulus onset there was a

suppression of alpha power contralateral to the winning sample. Although later in time than the modulation we found in broadband signals, previous studies have found similar ‘late’ suppression of alpha. Second, the timing, location and pattern of the parametric modulation resemble the findings of a previous study that measured fluctuations in human perceptual choice under focused and divided attention (Wyart et al., 2015). In that experiment, participants integrated tilt information from two simultaneous streams of Gabor patches under focussed attention (where the decision-relevant stream was signalled in advance) and divided attention (where it was not). In the focused condition, posterior electrodes contralateral to the attended stream, but not the unattended stream, showed a negative modulation around 250 ms. The reduced encoding of the losing sample in our study could similarly point to a focussing of attention towards the winning sample. Here we additionally show that this putative neural correlate of attention can be shifted rapidly between the two sides of the screen as samples arrive. Relatedly, the pattern of a posterior negativity contralateral to the target (in this case winning sample) shows a remarkable resemblance to the N2pc component in ERP literature. This component has long been related to (covert) attention and shifts in attention (Luck, 2012; Luck & Hillyard, 1994). Finally, when assessing how eye movements were influenced by decision information, we found that the position of the winning sample did shift the eyes starting from around 300 ms, despite our control for (larger) eye movements (see Figure D.2). This suggests a possible role of overt attention in selective integration alongside covert attention, although it is worth noting that eye movements only occurred after the observed posterior modulation of neural signals.

While our neural findings support selective integration as described by Tsetsos et al. (2012), they are also in line with other models making similar

theoretical claims. For example, the Associations and Accumulation Model (AAM) (Bhatia, 2013) similarly predicts that attention should be drawn to options where the attribute value is high, which in turn asymmetrically drives choices. The SI model also shares a lot of features with the Multi-alternative Decision by Sampling (MDbS) model (Noguchi & Stewart, 2018). One characterizing feature of the MDbS model is that attribute values are integrated through ordinal pairwise comparison, simply tallying attribute ranks. This is in contrast to the SI model, where the magnitude of this difference influences the strength of the evidence weighting. Both our model fits and the neural results suggest that a comparison between attribute values is not solely a process of ordinal comparison. We did find a (small) additive effect by rank (winner vs loser), encoded both in broadband and time-frequency signals. However, there was a much stronger parametric modulation between winning and losing samples that depended on the magnitude difference of the samples, which supports an additional multiplicative step not posited by MDbS.

More broadly, the models mentioned above set themselves apart from another class of sequential sampling models in their assumption that biased processing and integration occurs over a comparison of attribute values rather than the overall value of different options. The latter encompasses models such as the aDDM (Krajbich, Armel, & Rangel, 2010) and the MIVAC model (Gluth et al., 2018), which assume attention to an option boosts that option’s overall (expected) value. Because these studies used one stimulus that incorporated multiple attributes per option, a formal test of the SI model to this data is not possible. Future studies should be carried out that test whether the SI model is able to explain choices during simultaneous presentation of options with the possibility of sampling different attributes separately. Promisingly a recent study examined the contribution of a large

range of possible model parameters to predicting choice biases and found that a pairwise comparison of attributes, subjective weighting of attribute values and ‘leakage’ of information contribute significantly to improving model fits (Turner, Schley, Muller, & Tsetsos, 2018). The selective integration (SI) model incorporates all of these features.

In conclusion, our study provides a first step in understanding the neural mechanisms which accompany parallel integration and comparison of discrete samples of information.

Data and code availability

This work has been published in *Cerebral Cortex*: <https://doi.org/10.1093/cercor/bhaa039>. All data necessary to reproduce the results are available on <https://osf.io/cksj7/> (DOI 10.17605/OSF.IO/CKSJ7). All code and materials to reproduce the analyses and experiment are available at https://github.com/summerfieldlab/Luyckx_etal_2020_CerCortex.

References

- Bacigalupo, F., & Luck, S. J. (2019). Lateralized suppression of alpha-band EEG activity as a mechanism of target processing. *Journal of Neuroscience*, *39*(5), 900–917.
- Benjamini, Y., & Hochberg, Y. (2009). Controlling the False Discovery Rate : A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, *57*(1), 289–300.
- Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review*, *120*(3), 522–543.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, (10), 433–436.
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21.
- Donner, T. H., Siegel, M., Fries, P., & Engel, A. K. (2009). Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. *Current Biology*, *19*(18), 1581–1585.
- Glickman, M., Tsetsos, K., & Usher, M. (2018). Attentional selection mediates framing and risk-bias effects. *Psychological Science*, *29*(12), 2010–2019.
- Gluth, S., Spektor, M. S., & Rieskamp, J. (2018). Value-based attentional capture affects multi-alternative decision making. *eLife*, 1–36.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574.
- Hanks, T. D., & Summerfield, C. (2017). Perceptual decision making in rodents, monkeys, and humans. *Neuron*, *93*(1), 15–31.
- Kleiner, M., Brainard, D., & Pelli, D. (2007). "What's new in Psychtoolbox-3?" *Perception*, *36*(14).

- Krajovich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, *13*(10), 1292–1298.
- Li, V., Castañón, S. H., Solomon, J. A., Vandormael, H., & Summerfield, C. (2017). Robust averaging protects decisions from noise in neural computations. *PLoS Computational Biology*, *13*(8), e1005723.
- Luck, S. J. (2012). Electrophysiological correlates of the focusing of attention within complex visual scenes: N2pc and related ERP components. In *The oxford handbook of event-related potential components* (Chap. 12, pp. 329–360).
- Luck, S. J., & Hillyard, S. A. (1994). Electrophysiological correlates of feature analysis during visual search. *31*(3), 291–308.
- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, *164*(1), 177–190.
- Moran, R., & Tsetsos, K. (2018). The standard Bayesian model is normatively invalid for biological brains. *Behavioral and Brain Sciences*, *41*.
- Noguchi, T., & Stewart, N. (2018). Multialternative decision by sampling: A model of decision making constrained by process data. *Psychological Review*, *125*(4), 512–544.
- O’Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, *15*(12), 1–7.
- Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J. M. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*, *2011*(1-9).

- Pisauro, M. A., Fouragnan, E., Retzler, C., & Philiastides, M. G. (2017). Neural correlates of evidence accumulation during value-based decisions revealed via simultaneous EEG-fMRI. *Nature Communications*, 8, 15808.
- Sauseng, P., Klimesch, W., Stadler, W., Schabus, M., Doppelmayr, M., Hanslmayr, S., ... Birbaumer, N. (2005). A shift of visual spatial attention is selectively associated with human EEG alpha activity. *European Journal of Neuroscience*, 22(11), 2917–2926.
- Smith, P. L., & Little, D. R. (2018). In defense of the small-N design. *Psychonomic Bulletin and Review*, 25, 2083–2101.
- Spitzer, B., Waschke, L., & Summerfield, C. (2017). Selective overweighting of larger magnitudes during noisy numerical comparison. *Nature Human Behaviour*, 1, 1–8.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Group Bayesian Model Selection for Group Studies. *NeuroImage*, 46(4), 1004–1017.
- Summerfield, C., & Tsetsos, K. (2015). Do humans make good decisions? *Trends in Cognitive Sciences*, 19(1), 27–34.
- Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, 109, 9659–9664.
- Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences*, 113(11), 3102–3107.
- Turner, B. M., Schley, D. R., Muller, C., & Tsetsos, K. (2018). Competing theories of multialternative, multiattribute preferential choice. *Psychological Review*, 125(3), 329–362.

- Twomey, D. M., Murphy, P. R., Kelly, S. P., & Connell, R. G. O. (2015). The classic P300 encodes a build-to-threshold decision variable. *European Journal of Neuroscience*, *42*, 1636–1643.
- van Vugt, M. K., Simen, P., Nystrom, L. E., Holmes, P., & Cohen, J. D. (2012). EEG oscillations reveal neural correlates of evidence accumulation. *Frontiers in Neuroscience*, *6*, 1–13.
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior* (J. Von Neumann & O. Morgenstern, Eds.). Princeton: Princeton UP.
- von Lautz, A., Herding, J., & Blankenburg, F. (2019). Neuronal signatures of a random-dot motion comparison task. *NeuroImage*, *193*, 57–66.
- Wald, A., & Wolfowitz, J. (1949). Bayes solutions to sequential decision problems. *Proceedings of the National Academy of Sciences USA*, *35*(2), 99–102.
- Wyart, V., Myers, N. E., & Summerfield, C. (2015). Neural mechanisms of human perceptual choice under focused and divided attention. *Journal of Neuroscience*, *35*(8), 3485–3498.

Chapter 6

General Discussion

After four empirical chapters it may have become difficult for the reader to still see the forest through the trees. For this final chapter I would like to take a step back and take a look at the bigger picture. Rather than give a plain overview of our findings, I want to link them back to the models covered in Chapter 1 and to a new theory of posterior parietal cortex (PPC) function. In the final part of the chapter I will briefly review how our findings fit within the idea of adaptive representation of structure.

6.1 Structure learning models revisited

In Chapter 1, we covered different computational models of structure learning. We saw that the impact of these models on our understanding of the brain has by and large been limited, despite recent calls by for example Hassabis, Kumaran, Summerfield, and Botvinick (2017) for a closer interaction between neuroscience and computational modelling. Because this thesis has mainly focused on the neural representation of structure and the question of optimal representation, I will try to link my findings with the broader literature and critically review future directions.

First, **probabilistic models** assume that there exists a hypothesis space of possible structural forms. As we saw in Chapter 2, the human brain represented newly learned items with the same neural signals that were used for the well-known representation of numbers. It is possible that the signals for a one-dimensional structure were part of a larger hypothesis space and that the final structural representation was a consequence of the most plausible structural form. However, this remains speculation and future research is needed to test this hypothesis.

Furthermore probabilistic models indirectly deal with the question of optimality given their Bayesian approach. Moran and Tsetsos (2018) on the other hand suggest that Bayesian models lack integration or 'late' noise. Both Chapters 4 and 5 dealt with this question and showed that sequential sampling models allowing for non-linear transformation of the input guard against the effects of late noise. In Chapter 4 model simulations showed that anti-compressive transformations achieved higher accuracy than linear integration under late noise. Additionally normalizing within each number range led to greater robustness against late noise. Participants consistently exhibited the predicted anti-compressive transformation with a relative coding of the number space. In Chapter 5 these findings were extended to a scenario where concurrent information was being assessed. Previous theoretical work had already demonstrated the benefits of a transduction of the information based on its relative value compared to concurrent information (called 'selective integration'). We first showed that a model including a selectivity parameter was better at explaining the human choice behaviour. Importantly we observed an early difference in posterior regions in the encoding of items based on their relative value, with a stronger encoding of 'winning' samples. Finally, the model fits showed that participants with greater estimated noise

compensated with greater levels of transduction and this was not due to a trade-off between parameters in the model. Taken together these findings suggest that integration noise is an important factor in human information processing. While probabilistic models deal with the question of learning rather than acting upon input, it is worth considering what key features of human processing are still missing in these models.

Connectionist models start from a more low-level architecture, by updating the weights that govern the propagation of activity between layers of undifferentiated units. The simulations from Chapter 2 provided a proof-of-concept that neural networks are capable of representing structure without explicit feedback or with an adjusted architecture. While our network was successful at encoding and utilizing the latent structure of the inputs, its scope is very limited. The network was shallow (a single hidden layer) and we only tested near transfer to highly similar inputs. Promising work has come from Eslami et al. (2016) who built a recurrent neural network (RNN) that was able to identify multiple 2D or 3D objects in a scene through unsupervised learning. One interesting feature is that the model shifts attention across the scene to disentangle objects. In the next section, we will see that such a model should be well-suited to learn abstract structure representation according to Summerfield, Luyckx, and Sheahan (2020).

Other papers did look at the representations in deep neural networks (DNNs) and have found that number filters spontaneously arise in DCNNs trained for object recognition (DeWind, 2019; Nasr, Viswanathan, & Nieder, 2019). Nasr et al. (2019) describe how they first trained a DCNN to categorize objects from a large image database (Krizhevsky, Sutskever, & Hinton, 2012). Next, the network was presented with random dot arrays that varied in numerosity, size and spacing. They found that in the final layer of the

network a small portion of units responded to the numerosity of the array independent of size or spacing. DeWind (2019) further found units in earlier layers of a DCNN that activated linearly with numerosity similar to number-sensitive neurons in PPC (Roitman, Brannon, & Platt, 2007) (see below). One of the limitations translating these findings from DNNs to human structure learning is that humans can be successful at these tasks without having access to a vast amount of labelled data. But it does emphasize an important role of visual processing in structure learning. It also speaks against the idea of an innate system of number sensitivity in favour of the idea that it arises through interaction with the environment, in line with the 'emergentist' view. It would be interesting to see whether the magnitude representation found in DNNs can also translate to new tasks as in Chapter 2, a key feature of true structure representation.

Finally, **reinforcement learning models** use a state space of interconnected states as an abstract structure representation. Moreover a hierarchical representation of the state space allows these models for more flexible decision-making and faster learning. Perhaps the findings most relevant to HRL come from Chapter 3. There we provided direct evidence that magnitude was represented neurally along a single axis, but separated by context. Participants benefited from actively representing the context, because it allowed them to make educated guesses about the correct answer in case they forgot the previous number or missed the current number. Encoding of distance to context mean was both found in behavioural and univariate EEG measurements.

We further saw that the overlap between spatial and conceptual navigation has been an important factor in understanding the neural processes linked to reinforcement learning, mainly in identifying the medial temporal lobe as a key region. As we will see in the following section, the results from this thesis

could point towards an analogue system to MTL in humans in the posterior parietal cortex (PPC).

6.2 Posterior parietal cortex as candidate region

Due to the low spatial resolution of EEG, this thesis has so far shied away from any claims about the involvement of specific brain regions. However, several of the findings presented in previous chapters fit within a new theory about the function of posterior parietal cortex (PPC) (Summerfield et al., 2020). According to this new theory, primate PPC encodes low-dimensional abstract representations of structure. This would be in line with suggestions that for the brain to generalize structure to new contexts, it needs to orthogonalize structure from content (Behrens et al., 2018). Within this view, the ventral stream processes the sensory content of a scene, while the dorsal stream processes the relational structure within a scene. The theory draws parallels with the MTL in rodents, a region that seems important for both spatial and conceptual navigation. The close link between the two sorts of navigation may indicate that structure representation is learned through an initial understanding of how space is organized. And while rodents navigate their environment allocentrically, primates more often work from a peripersonal frame of reference. The theory thus suggests that the PPC, through its close link with the saccadic system, plays an important role in scene understanding in primates and consequently in structure understanding.

The theory builds on 'A Theory Of Magnitude' (ATOM) (Buetti & Walsh, 2009; Walsh, 2003) which argues that parietal cortex contains a common neural code for magnitude that abstracts over concepts of space, time and number. PPC has long been implicated in numerical cognition (Eger et al., 2009; Nieder, 2016; Nieder & Dehaene, 2009; M. Piazza, Pinel, Le Bihan, &

Dehaene, 2007; Manuela Piazza, Izard, Pinel, Le Bihan, & Dehaene, 2004). Roitman et al. (2007) for example found number-sensitive neurons that monotonically increased (or decreased) activity with increasing numerical magnitude, as if magnitude lay on a single line. Viswanathan and Nieder (2013) found number-selective neurons in primate PPC that fired with Gaussian tuning curves around a preferred number (although Chen and Verguts (2013) have argued that a number-sensitive code could still be underlying these findings). Interestingly, overlapping regions in PPC have also been found to code for space (Chafee, Averbeck, & Crowe, 2007; Oliveri et al., 2009) and/or time (Bonato, Zorzi, & Umiltà, 2012; Jazayeri & Shadlen, 2015; Oliveri et al., 2009).

PPC activity has been found in a much wider range of studies involving for example evidence strength (de Lafuente, Jazayeri, & Shadlen, 2015; Shadlen & Newsome, 2001), value (Platt & Glimcher, 1999) or categorization (Fitzgerald et al., 2013) (Figure 6.1). According to the new theory, these findings can be reinterpreted within the view that PPC projects high-dimensional information into a low-dimensional manifold. For example, Rutishauser, Aflalo, Rosario, Pouratian, and Andersen (2018) found that activity of single neurons in human PPC scaled with memory strength and confidence (Figure 6.2A). Cao, Summerfield, Park, Giordano, and Kayser (2019) in a multisensory experiment showed that intraparietal sulcus (IPS), in contrast to primary auditory cortex, integrated visual and auditory rate to represent it along a single axis (Figure 6.2B). Alfred, Connolly, and Kraemer (2018) and Alfred, Connolly, Cetron, and Kraemer (2020) have shown that PPC is involved during transitive inference, a hallmark of a one-dimensional representation where items are positioned on a line.

We can interpret the findings from Chapters 2 and 3 within the same




Figure 6.1. Examples of scalar codes in primate parietal cortex. Extension of ATOM theory (Buetti & Walsh, 2009; Walsh, 2003) which argues that overlapping neural codes for space, time and number in the parietal cortex are all manifestations of a common representation of magnitude. Summerfield, Luyckx, and Sheahan (2020) suggest that other quantities that are coded monotonically in PPC cells - including value, decision information, and even arbitrary categories - may reflect the projection of high-dimensional signals onto a one-dimensional (or low-dimensional) manifold that is suited for generalisation and transfer, i.e. for the understanding of the analogical relations between abstract concepts. Here, we emphasise those concepts that share one-dimensional structure.

framework. First, univariate activity in centro-parietal (Chapter 2) or right parietal electrodes (Chapter 3) varied strongly with the numerical magnitude of the stimuli (Figure 6.2C). Note that numbers were in symbolic format, which excludes differences in visual stimulation as an alternative explanation. In Chapter 2 we further investigated the number of dimensions needed to successfully describe the cross-validation between the numerical and value-based task. By systematically reconstructing the dimensionality of the data and assessing the strength of cross-validation between the two tasks, we found that there was no significant improvement in cross-validation after adding just two dimensions. MDS further confirmed that a two-dimensional space was sufficient to visualize the configuration of stimuli along a common magnitude axis (Figure 6.2E). In Chapter 3 we used a dimensionality reduction technique to visualize the first 3 dimensions of the neural signal. We found that the representation of magnitude could be projected onto a single axis in the third principal dimension.

For completeness, some authors have questioned the claim that there exists a generalized magnitude system in parietal cortex (Borghesani et al., 2019; Cavdaroglu & Knops, 2018; Van Opstal & Verguts, 2013). Borghesani et al. (2019) presented participants with a variant of the experiment in Chapter 2. Participants in two conditions compared either the average number of dots or the average length of lines between two sets. Using RSA they found that right parietal cortex encoded both number and line length, but they could not find any evidence that the region also encoded a domain-general magnitude metric. According to Van Opstal and Verguts (2013) an alternative explanation for our results could be that the shared magnitude representation arises at the level of working memory rather than as the domain-general coding of a brain region. They argue that the overlap is only found when




Figure 6.2. Low-dimensional coding of magnitude in human PPC. **(A)** Example single neuron in human PPC whose firing rate depends on memory strength as reported on a numerical scale, with lowest response for number '6' and highest for number '1' (Rutishauser et al., 2018). **(B)** Whilst undergoing MEG recordings, humans judged the rate of a multisensory stimulus that pulsed simultaneously and independently in the visual (flashes) and auditory (flutter) domains. Left panel: colours show different auditory and visual rate conditions. Right panel: multidimensional scaling (MDS) reveals that brain signals over PPC represented fused multisensory rate with a low-dimensional signal (Cao et al., 2019), whereas A1 coded for the input rate. **(C)** Reiteration of the findings from Chapter 2. Univariate EEG signals evoked by symbolic numbers (1-6; top panel) and visual images that paid out reward with fixed probability (ranked by their value, b1-b6; bottom panel). **(D)** Cross-validated RSA shows that common coding for numbers and bandits with the same magnitude occurred at 300-700 ms post-stimulus. **(E)** MDS shows the low-dimensional form of the shared neural representation.

magnitude is relevant for task context, indicating the need to be actively encoded. In Chapter 3 we did not find any representation of magnitude for filler numbers, when no magnitude judgement was required, which would be in line with this explanation. Unfortunately the spatial resolution of EEG in our experiments does not allow us to test these hypotheses in more detail. Neither do the domain-general and working memory hypotheses have to be mutually exclusive. One possibility is that there exist regions in the brain that contain invariant representations of structure, while other regions can build, combine and transform structures to fit the task at hand. The Hierarchical Bayesian models described in Chapter 1 demonstrate the feasibility and flexibility of such a hierarchically structured system, with lower levels encoding basic forms that can be combined into complex structures (Lake, Salakhutdinov, & Tenenbaum, 2015; Tervo, Tenenbaum, & Gershman, 2016).

6.3 Flexible adaptive structure representation?

As a common thread throughout all empirical chapters, a second question has been whether structure representation - and information-processing in general - can flexibly adapt to the requirements of the task. In Chapter 2 we replicated the finding from Spitzer, Waschke, and Summerfield (2017) that (a) decision weights exponentially increased with numerical magnitude (as opposed to linearly) and (b) the level of anti-compression correlated strongly with the distortion estimated from neural patterns. Previous work already established that such anti-compression leads to better performance when one assumes accumulation or 'late' noise. We found the same distortion in decision-weights in Chapter 4, which used an almost identical task. Moreover a model that employed a single compression parameter over different numerical ranges best explained the choice data. Whether these results are specific to the

numerical integration task is hard to say. Although not included in this thesis, attempts to estimate the level of distortion in neural signals for the numerical comparison task in Chapter 3 were not successful.

Normalization of the input is another feature of efficient coding that protects information-processing against late noise. In Chapter 3, our results indicated that the number lines were normalized within their respective ranges, both in the neural and in the behavioural data. In Chapter 4 we explored the implications of normalization in more detail. Model simulations indicated that normalizing with each number range led to greater robustness when faced with late noise. In line with the theoretical predictions, we found that the best fitting model to the human choices normalized within each range.

Finally, in Chapter 5 we found neural evidence for a model that performed a non-linear transformation of incoming information before accumulation, similar to the model used in Chapter 2 and 4. While the first three empirical chapters used a task with a single stream of stimuli, Chapter 5 extended the case to a dual-stream of information. Accordingly the biased weighting of information no longer applied to individual items, but to the relative evidence that an item carried within an item pair. Overall, these results point towards a neural system that optimizes its performance given the noise inherent in the system. More work is needed to understand how flexible structure representation and distortion extend to other paradigms beyond sequential sampling. Future work can also focus on how the distortions may extend to cases with more than two items present.

6.4 Concluding remarks

The observant reader may have noticed that while I introduced a whole range of structures at the beginning of this thesis, the focus has remained on the one-dimensional representation of numerical magnitude. While this is an entirely fair point, I will try to substantiate our approach. First, to understand a complex problem such as structure learning it is instructive to start with the most basic case. Second, the studies were motivated by theoretical work that has suggested 'number' is one expression of an abstract magnitude representation that also applies to space and time perception (Bueti & Walsh, 2009; Fischer & Shaki, 2018; Sokolowski, Fias, Bosah Ononye, & Ansari, 2017; Walsh, 2003). Finally, the work by Spitzer et al. (2017) (and reiterated in Chapter 2) should have made it clear that the representation of magnitude we observed in EEG signals was independent of both task and sensory information.

The contribution of this thesis does not lie in its insights in the processes of learning complex structures, but rather in how the brain comes to represent those structures and then utilizes that representation. It is up to future studies to determine how our findings translate to different or more complex structures. The structure learning models can deal with complex structures, but more work is needed to link the mechanisms suggested in these models to the neural mechanisms in the brain.

This thesis has only scratched the surface of structure representation and learning. I hope that it can serve as a starting point to further deepen our understanding of the neural processes linked to structure representation. And while it has focused on neural processes, it should also serve as a reminder that as psychologists we should not lose sight of the insights we have gained

from human behaviour. Only then can we build models that truly reflect the human nature.

References

- Alfred, K. L., Connolly, A. C., Cetron, J. S., & Kraemer, D. J. (2020). Mental models use common neural spatial structure for spatial and abstract content. *Communications Biology*, 3(1), 1–11.
- Alfred, K. L., Connolly, A. C., & Kraemer, D. J. (2018). Putting the pieces together: Generating a novel representational space through deductive reasoning. *NeuroImage*, 183, 99–111.
- Behrens, T. E., Muller, T. H., Whittington, J. C., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron*, 100(2), 490–509.
- Bonato, M., Zorzi, M., & Umiltà, C. (2012). When time is space: Evidence for a mental time line. *Neuroscience and Biobehavioral Reviews*, 36(10), 2257–2273.
- Borghesani, V., de Hevia, M. D., Viarouge, A., Pinheiro-Chagas, P., Eger, E., & Piazza, M. (2019). Processing number and length in the parietal cortex: Sharing resources, not a common code. *Cortex*, 114, 17–27.
- Bueti, D., & Walsh, V. (2009). The parietal cortex and the representation of time, space, number and other magnitudes. *Philosophical Transactions of the Royal Society B*, 364(1525), 1831–1840.
- Cao, Y., Summerfield, C., Park, H., Giordano, B. L., & Kayser, C. (2019). Causal Inference in the Multisensory Brain. *Neuron*, 102(5), 1076–1087.e8.
- Cavdaroglu, S., & Knops, A. (2018). Evidence for a Posterior Parietal Cortex Contribution to Spatial but not Temporal Numerosity Perception. *Cerebral Cortex*, (July), 1–13.

- Chafee, M. V., Averbeck, B. B., & Crowe, D. A. (2007). Representing spatial relationships in posterior parietal cortex: Single neurons code object-referenced position. *Cerebral Cortex*, 17(12), 2914–2932.
- Chen, Q., & Verguts, T. (2013). Spontaneous summation or numerosity-selective coding? *Frontiers in Human Neuroscience*, 7(DEC), 1–3.
- de Lafuente, V., Jazayeri, M., & Shadlen, M. N. (2015). Representation of accumulating evidence for a decision in two parietal areas. *Journal of Neuroscience*, 35(10), 4306–4318.
- DeWind, N. K. (2019). The number sense is an emergent property of a deep convolutional neural network trained for object recognition. *bioRxiv*, 1–31.
- Eger, E., Michel, V., Thirion, B., Amadon, A., Dehaene, S., & Kleinschmidt, A. (2009). Deciphering Cortical Number Coding from Human Brain Activity Patterns. *Current Biology*, 19(19), 1608–1615.
- Eslami, S. M. A., Heess, N., Weber, T., Tassa, Y., Szepesvari, D., Kavukcuoglu, K., & Hinton, G. E. (2016). Attend, infer, repeat: Fast scene understanding with generative models. *Advances in Neural Information Processing Systems*, 3225–3233.
- Fischer, M., & Shaki, S. (2018). Number concepts: Abstract and embodied. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752), 20170125.
- Fitzgerald, J. K., Freedman, D. J., Fanini, A., Bennur, S., Gold, J. I., & Assad, J. A. (2013). Biased associative representations in parietal cortex. *Neuron*, 77(1), 180–191.
- Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron*, 95(2), 245–258.

- Jazayeri, M., & Shadlen, M. N. [Michael N.]. (2015). A neural mechanism for sensing and reproducing a time interval. *Current Biology*, *25*(20), 2599–2609.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Handbook of approximation algorithms and metaheuristics. In *Advances in neural information processing systems* (pp. 1097–1105).
- Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, *350*(6266), 1332–1338.
- Moran, R., & Tsetsos, K. (2018). The standard Bayesian model is normatively invalid for biological brains. *Behavioral and Brain Sciences*, *41*.
- Nasr, K., Viswanathan, P., & Nieder, A. (2019). Number detectors spontaneously emerge in a deep neural network designed for visual object recognition. *Science Advances*, *5*(5), 1–10.
- Nieder, A. (2016). The neuronal code for number. *Nature Reviews Neuroscience*, *17*(6), 366–382.
- Nieder, A., & Dehaene, S. (2009). Representation of number in the brain. *Annual Review of Neuroscience*, *32*(1), 185–208.
- Oliveri, M., Koch, G., Salerno, S., Torriero, S., Gerfo, E. L., & Caltagirone, C. (2009). Representation of time intervals in the right posterior parietal cortex: Implications for a mental time line. *NeuroImage*, *46*(4), 1173–1179.
- Piazza, M. [M.], Pinel, P., Le Bihan, D., & Dehaene, S. (2007). A magnitude code common to numerosities and number symbols in human intraparietal cortex. *Neuron*, *53*(2), 293–305.

- Piazza, M. [Manuela], Izard, V., Pinel, P., Le Bihan, D., & Dehaene, S. (2004). Tuning curves for approximate numerosity in the human intraparietal sulcus. *Neuron*, 44(3), 547–555.
- Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, 400, 233–238.
- Roitman, J. D., Brannon, E. M., & Platt, M. L. (2007). Monotonic coding of numerosity in macaque lateral intraparietal area. *PLoS Biology*, 5(8), 1672–1682.
- Rutishauser, U., Aflalo, T., Rosario, E. R., Pouratian, N., & Andersen, R. A. (2018). Single-Neuron Representation of Memory Strength and Recognition Confidence in Left Human Posterior Parietal Cortex. *Neuron*, 97(1), 209–220.e3.
- Shadlen, M. N. [M. N.], & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86, 1916–1936.
- Sokolowski, H. M., Fias, W., Bosah Ononye, C., & Ansari, D. (2017). Are numbers grounded in a general magnitude processing system? A functional neuroimaging meta-analysis. *Neuropsychologia*, 105, 50–69.
- Spitzer, B., Waschke, L., & Summerfield, C. (2017). Selective overweighting of larger magnitudes during noisy numerical comparison. *Nature Human Behaviour*, 1, 1–8.
- Summerfield, C., Luyckx, F., & Sheahan, H. (2020). Structure learning and the posterior parietal cortex. *Progress in Neurobiology*, 184, 101717.
- Tervo, D. G. R., Tenenbaum, J. B., & Gershman, S. J. (2016). Toward the neural implementation of structure learning. *Current Opinion in Neurobiology*, 37, 99–105.

- Van Opstal, F., & Verguts, T. (2013). Is there a generalized magnitude system in the brain? Behavioral, neuroimaging, and computational evidence. *Frontiers in Psychology*, 4(JUL), 2011–2013.
- Viswanathan, P., & Nieder, A. (2013). Neuronal correlates of a visual "sense of number" in primate parietal and prefrontal cortices. *Proceedings of the National Academy of Sciences*, 110(27), 11187–11192.
- Walsh, V. (2003). A theory of magnitude: Common cortical metrics of time, space and quantity. *Trends in Cognitive Sciences*, 7(11), 483–488.

Appendix A

Supplementary material to Chapter 2




Figure A.1. RSA results with Pearson correlation distance. We repeated all main analyses using correlation distance, a measure that is scale invariant and removes any univariate effects, showing that our results are consistent over different distance measures. **(A)** RSA revealed a magnitude effect arising around 100 ms for both tasks (bottom coloured lines, $P_{cluster} < 0.005$). Shaded area represents SEM. **(B)** Averaged RDMs between 200 – 700 ms for both tasks. **(C)** Averaged neural RDM between 350-600 ms across the two tasks. Upper left and lower right quadrant represent within-task dissimilarities for the numerical and bandit task respectively, lower right and upper left quadrant contain the cross-validated dissimilarities between tasks. **(D)** Cross-temporal cross-validation revealed a stable magnitude representation shared over the two tasks between approximately 350 – 600 ms ($P_{cluster} < 0.005$).




Figure A.2. RSA results split per task framing. Approximately half of the participants performed the numerical task indicating the color with the lowest average ($N = 24$) and the other half indicating the highest average ($N = 22$). **(A-B)** Within-task RSA showed magnitude representations were reliably decodable for the two tasks under either task framing, albeit with slightly later/smaller clusters than the aggregated data (bottom coloured lines, $P_{cluster} < 0.005$). Shaded area represents SEM. **(C-D)** Cross-temporal cross-validation showed a shared magnitude representation under either task framing, with a slightly larger cluster for the higher frame ($P_{cluster} < 0.005$).




Figure A.3. RSA results split for first and second bandit. To test for the possibility that our magnitude decoding was not merely driven by a ‘value comparison’ signal, whereby the value of the second bandit is assessed relative to the value of the first, RDMs were calculated separately for the first and second presented bandit. **(A)** Although the identified cluster was greater for the second bandit, magnitude was decodable for both bandits reliably from approximately 350 ms onwards (bottom coloured lines, $P_{cluster} < 0.005$). **(B)** Averaged neural RDMs between 200 – 700 ms for the first (top) and second (bottom) bandit were highly similar.




Figure A.4. Cross-validation after removal of number 6/subjectively highest valued bandit (subset of Figure 2.2 C). The most distinct pattern of numerical/value distance was present for the highest number and most valued bandit. After removing this pair from the cross-validated RDM and repeating the cross-temporal cross-validation RSA, magnitude was still decodable from the remaining numbers and bandits ($P_{cluster} < 0.005$).




Figure A.5. We compared the strength of within-subject cross-validation to the between-participants cross-validation, asking whether variations in multivariate neural signals existed that were idiosyncratic within participants. For each participant, we combined EEG responses per stimulus from the numerical task iteratively with the EEG responses of another participant from the bandit task, repeating our cross-validation protocol for the newly created between-subject data. This allowed us to ask how well a participant’s neural patterns correlated with those of other participants, resulting in a distribution of correlation estimates at each time point. We then calculated a z-value at each time point, indicating how far the within-participant correlation diverged from the between-participant distribution of correlations. Non-parametric cluster-based permutation tests were then used to identify clusters in time where participants’ neural patterns in the bandit task were significantly better predicted by their own neural patterns compared to the population ($P_{cluster} < 0.005$). The effect was noisier but arose around the same time as the original cross-validation pattern, suggesting the effect was not purely driven by shared noise from the same recording session.




Figure A.6. Univariate centro-parietal signals separated by task and task framing. Grey shaded areas show time of greatest disparity between signals (Kruskal-Wallis, $P_{FDR} < 0.01$). **(A-B)** Centro-parietal responses for numbers followed an ascending pattern from low to high numbers, independent of whether the task instructions asked to indicate the colour with the lowest or highest average. Scalp maps show response amplitude for number 6 at the identified time window. **(C-D)** Centro-parietal responses for bandits followed the same ascending pattern, going from subjectively lowest valued bandit to the highest valued bandit. The bandit task was identical for both framing groups. Scalp maps show response amplitude for the highest valued bandit b6 at the identified time window.




Figure A.7. We tested whether the univariate CPP amplitude modulations could explain our multivariate findings. The magnitude model was compared to a model RDM constructed from univariate CPP amplitude differences in a multiple regression. **(A-B)** CPP amplitudes were obtained by averaging the univariate centro-parietal response to each stimulus (digit or bandit) within the individually identified time windows from Figure 2.5 A-B. CPP RDMs (inset) were constructed based on the differences between average amplitudes. Error bars represent SEM. **(C)** Average Euclidean distances of CPP amplitudes from both tasks combined (numerical: 1-6, bandit: b1-b6). The lower left quadrant, containing the between-task distances in CPP amplitudes, served as the CPP model RDM for the cross-validation analysis. **(D-F)** For each control analysis, the magnitude model significantly explained variance in the multivariate neural patterns even after entering the univariate CPP as a second regressor in a multiple regression (within-task comparison: $P_{cluster} < 0.005$, bottom coloured lines; cross-validation: $P_{cluster} < 0.05$). Shaded coloured area represents SEM.




Figure A.8. Cross-validation after excluding the first principal dimension. A more stringent test to control for univariate effects is to exclude the first dimension identified by SVD. Cross-temporal cross-validation on the reduced data revealed a smaller but significant cluster at the same time points as our main result and a previously unobserved later cluster ($P_{cluster} < 0.05$). Consistent with our previous analyses (Figure 2.5 A-C), this suggests there is a major univariate component to our multivariate findings, but the shared pattern also exists in higher dimensions of the data.

Appendix B

Supplementary material to Chapter 3




Figure B.1. RSA control analysis. To further assess whether context separation (red) was not explained by distinct colour mapping of conditions or magnitude (blue) by motor effects, we repeated the RSA with two additional, subject-specific model RDMs: (a) colour dissimilarity (yellow) and (b) reaction time (black). Colour similarity now explained early representational structure in the neural signals, while context representation explained the later time points from around 400 ms. Reaction time measures arose around the same time as magnitude representation, but notably did not explain away the variance explained by magnitude representation. Insets show the model RDMs (example subject for right panels). Shaded area indicates SEM. Lines below indicate significant clusters after cluster-correction ($P_{cluster} < 0.05$).




Figure B.2. Examples configuration model. To assess whether number lines were parallel on average, we designed a model that recreated a three-dimensional configuration of three number lines, similar to the MDS projection seen in Figure 3.4A. The angle and direction of each number vector could vary independently, which in a next step could be used to assess the degree of "parallelness" between lines by calculating the pairwise cosine distance between angles (Figure 3.4C). 3D configurations (top) were translated into RDMs (bottom) to allow fitting to the neural data. The model was able to recreate orthogonal configurations (left panel), vary the distance between the lines (second panel), alter the projection onto each axis (third panel) and even let number vectors point in opposite direction (right panel). Numbers were equally spaced along the direction of the vector from low to high.

Appendix C

Supplementary material to Chapter 4




Figure C.1. Model fits for mixed trials. We repeated the same model fitting procedure from Section 4.2.3 (p. 104) using only trials from blocks where both ranges were interleaved. Both log-likelihood estimates and protected exceedance probability favoured a model with relative scaling and a single compression parameter k . This suggests participants were able to flexibly switch between the two ranges and did not resort to placing numbers on a single number line.




Figure C.2. Assessing inflection point around decades. By using two-digit numbers, one alternative explanation for a non-linear shape of the number line is a break around the decades (30 and 40). In order to visualize whether decades served as a natural inflection point, we plotted the estimated decision weights, separated by range and block type. The shape of the number line was consistent over different shifts of the ranges, with no apparent influence of the decade. Note that estimates might be unreliable here because the number of participants in some ranges is relatively low ($5 < n < 10$).

Appendix D

Supplementary material to Chapter 5

Parameter recovery

To ensure there was no trade-off between the parameters of our model, we tested whether simulated parameter combinations could be recovered using the same model-fitting procedure as above. We were mainly interested in the possible trade-off between the weighting parameter m and late noise parameter s and therefore performed our parameter recovery procedure once for each parameter of interest.

First, we generated 25 equidistant values over the full sampling range of the parameter of interest. We then generated 500 random combinations of the three remaining parameters. The sampling space of all parameters was limited by the minimum and maximum estimated values from the real data. At every step of the parameter of interest, model responses were generated for the 500 parameter combinations using the input data from a randomly selected subject each time (1200 trials). Finally, the input data and model generated




Figure D.1. Parameter recovery. Parameter recovery error was not significantly different from zero (*all* $P > 0.05$), both when systematically varying m (left two panels) or s (right two panels). The larger confidence interval for κ when varying s likely indicates a trade-off between the slope and lapse rate of the response function. Error bars represent 95% confidence interval.

responses were used to estimate best-fitting parameters again and test how well we could recover the original parameter combinations. Statistical tests on the recovery error for the remaining three parameters were performed on the 25 values obtained after averaging over the 500 random iterations at each step.

EOG control analysis: regression model

Although we tried to control for large saccades with the eye-tracker, there remained the possibility that small eye movements towards the winning sample were driving some of our neural findings. Because we did not have access to very precise eye-tracking data as it was only used for online monitoring, we turned to the signals recorded with the horizontal dipolar electrodes (HEOG). To assess how eye movements were influenced by decision information, here operationalised as the difference in momentary evidence, we constructed a regression model with the epoched HEOG signal as the dependent variable.

$$HEOG_k = \beta_0 + \beta_1 \cdot |\Delta(X_k^R, X_k^L)| + \beta_2 \cdot \text{sgn}(\Delta(X_k^R, X_k^L)) + \varepsilon \quad (\text{D.1})$$

The difference in decision information ($\Delta(X_k^R, X_k^L)$) was split into two regressors: the absolute strength of evidence and the sign (or direction) of the information. The regression was then performed for each sample epoch independently at each time point and electrode, identical to the single-trial EEG regression.

Results

We found that indeed eye movements were significantly explained by the sign of the decision information between ~ 300 and 500 ms ($P_{\text{cluster}} < 0.05$; Figure D.2), but not by the strength of the evidence. In other words, the eyes would move towards the winning sample. Importantly, these effects only arose around 300 ms and peaked at ~ 400 ms, after the observed parametric modulation of decision evidence in contralateral electrodes (Figure 5.3, p. 140). This suggests that the modulation of posterior EEG signals is unlikely

to be driven by small shifts in eye movements. Nevertheless, we decided to include HEOG channel data as a nuisance regressor in subsequent analyses to exclude the possibility of eye movement contamination.




Figure D.2. Eye movement control analysis. When regressing HEOG channel data against the momentary evidence, the sign but not the strength of decision evidence, significantly predicted eye movements starting around 300 ms. Crucially, these effects arose later than our lateralized neural findings. Shaded coloured area indicates SEM.

Control analysis: GLM with model input

To further establish a link between our model fitted parameters and the observed modulation in EEG signals, we repeated our regression analysis (Equation 5.9), replacing the generative sample values by the model transformed sample values after selective weighting. We reasoned that if the difference in EEG signals for winning and losing samples was an indication of selective weighting, then this difference should disappear when entering the values in the model that were already transformed by the mode. Indeed, we observed the same early posterior negativity for both winning and losing samples around 200 - 300 ms, but there was no longer a significant difference in the strength of the modulation ($P_{cluster} > 0.05$). We do note that we did not observe a correlation between the model estimated m parameter and the difference in early posterior negativity across subjects in a cross-validated time window around 200 - 300 ms (Spearman $\rho = 0.35$, $P = 0.21$), possibly due to our small sample size.




Figure D.3. EEG regression analysis with model estimated values. When using the model transformed sample values as input to the regression, we no longer observed a significant difference between the modulation in EEG signals for winning and losing samples.

Individual model fitting results

ID	$\mathbf{m}$	λ	κ	$\mathbf{s}$	$\mathbf{LL}_{\text{fixed}}$	$\mathbf{LL}_{\text{full}}$
1	0.0004	0.1191	0.0000	0.0935	-526.54	-527.72
2	0.0572	0.2119	0.0386	0.0205	-658.74	-643.22
3	0.0509	0.1684	0.0675	0.0186	-698.88	-711.32
4	-0.0004	0.1384	0.0000	0.0909	-598.25	-599
5	0.0333	0.2902	0.0000	0.0343	-567.93	-575.55
6	0.0035	0.3063	0.0603	0.0968	-670.1	-671.84
7	0.1187	0.3260	0.0000	0.0164	-707.01	-664.21
8	0.0544	0.3134	0.0120	0.0198	-669.34	-661.61
9	0.0397	0.4236	0.0000	0.0242	-679.32	-673.31
10	0.0014	0.1171	0.0267	0.0756	-609.2	-611.42
11	0.0395	0.0827	0.0260	0.0235	-535.95	-526.81
12	0.0501	0.2124	0.0227	0.0286	-571.56	-562.29
13	0.0532	0.1652	0.0000	0.0181	-628.48	-628.32
14	0.0202	0.3191	0.0000	0.0319	-687.96	-687.57
15	0.0429	0.3155	0.0096	0.0316	-607.33	-609.48

Table D.1. Model fitting results. Individual parameter estimates from the full model and negative log-likelihood estimates for the cross-validated fixed and full model.