



Diaz, Jessica Ann (2018) The neural correlates of the influence of learning on perceptual decision making. PhD thesis.

<https://theses.gla.ac.uk/30795/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk/>
research-enlighten@glasgow.ac.uk



University
of Glasgow | Institute of Neuroscience
& Psychology

THE NEURAL CORRELATES OF THE
INFLUENCE OF LEARNING ON
PERCEPTUAL DECISION MAKING

JESSICA ANN DIAZ

LICENCIATURA IN PSYCHOLOGY; MSc BRAIN IMAGING

SUBMITTED IN FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE OF
Doctor of Philosophy

INSTITUTE OF NEUROSCIENCE & PSYCHOLOGY

COLLEGE OF MEDICAL, VETERINARY & LIFE SCIENCES

UNIVERSITY OF GLASGOW

OCTOBER, 2017

Abstract

Perceptual decision making involves the classification of sensory information, usually followed by an overt behavioural response. Any decision making, and perceptual decision making in particular, can be understood both theoretically and neurologically as a process of an accumulation of evidence to some threshold, at which point a commitment to a choice is made. This process can be examined in human subjects by analysing EEG data during perceptual decision making and identifying temporal components that are the neural signatures of the accumulation-to-bound decision making (see Philiastides & Sajda, 2006; Philiastides, Ratcliff, & Sajda, 2006a; Philiastides et al., 2006a; Ratcliff, Philiastides, & Sajda, 2009b). It may also be statistically modelled using sequential sampling models (see Ratcliff & Smith, 2004; Ratcliff, Gomez, & McKoon, 2004; Ratcliff & McKoon, 2008; Ratcliff & Van Dongen, 2011). Taken together, these provide us with a experimental and theoretical framework for the study of the neuroscience of human decision making. In this thesis, our aim is to address some open questions with respect to human perceptual decision making using the theoretical framework of sequential sampling models and the experimental paradigm of measuring temporal components in single-trial EEG discriminant analysis.

In Chapter 2, we will describe our studies of the role of learning on perceptual decision making. In particular, here we address competing hypotheses about the nature and location of perceptual learning in the brain. We provide evidence that perceptual learning arises from changes in higher level brain areas that are related to decision-making, rather than from perceptually earlier areas that are related to the encoding of sensory stimuli. In Chapter 3, we provide a specific mechanistic account of how learning affects perceptual decision making. This work follows from the work of others who have applied reinforcement learning theories (see Sutton & Barto, 1998) to the study of perceptual learning. In Chapter 4, we will describe our studies of the interaction of prior expectation and learning on decision making. In this study, we particularly aim to address whether prior expectation affects baseline activation or evidence accumulation in the decision making system, and how this changes with training. Here, we obtain evidence showing how the effects

of prior expectation are more related to evidence accumulation rather than baseline activation. In Chapter 5, we provide sequential sampling models, particularly drift diffusion models, of the data that we've obtained in the main experiments described in Chapter 2 and Chapter 4. The principal results here show how learning and prior expectation primarily have their effect on perceptual decision making by increasing the rate of evidence accumulation.

Our general conclusion is that using a combination of the theoretical framework of sequential sampling models and the experimental paradigm of measuring temporal components in single-trial EEG discriminant analysis provides an effective and comprehensive means to address open questions with respect to human perceptual decision making.

Contents

1	General introduction	1
1.0.1	Statistical decision theory	3
1.0.2	Sequential sampling models	4
1.0.3	EEG and the neural signatures of decision making	8
1.0.4	Interim summary	9
1.1	The effects of learning on perceptual decision making	10
1.1.1	“Early” perceptual learning	11
1.1.2	“Intermediate” perceptual learning	13
1.1.3	“Late” perceptual learning	14
1.1.4	Interim summary	15
1.2	The role of prior expectation on perceptual decision making	16
1.3	Aims and outline of the thesis	19
2	Perceptual learning alters post-sensory processing in human decision making	22
2.1	Introduction	22
2.2	Materials and Methods	25
2.2.1	Subjects	25
2.2.2	Stimuli	25
2.2.3	Behavioural Task	26
2.2.4	EEG Data Acquisition	27
2.2.5	Single Trial EEG Data Analysis	27
2.2.6	Multilevel Regression Analyses	32
2.3	Behavioural analysis results	36
2.4	EEG analysis results	40

2.4.1	Analysis of peak A_z values	40
2.4.2	Analysis of component onset times	43
2.4.3	Analysis of the discriminating component y	43
2.4.4	Prediction of behavioural responses by early versus late component	46
2.5	Discussion	49
3	A reinforcement learning account of the effect of learning on perceptual decision making	53
3.1	Reinforcement learning accounts of perceptual learning	53
3.1.1	Extension by Kahnt et al	56
3.2	Modelling face versus car perceptual learning	57
3.3	Results of modelling face versus car perceptual learning	59
3.3.1	Analysis of the y variable	59
3.4	Extensions	62
3.5	Discussion	65
4	The neural signatures of the role of prior expectation in perceptual decision making	68
4.1	Introduction	68
4.2	Materials and Methods	70
4.2.1	Subjects	70
4.2.2	Stimuli	71
4.2.3	Behavioural Task	71
4.2.4	EEG Data Acquisition	73
4.2.5	Single Trial EEG Data Analysis	73
4.3	Results	73
4.3.1	Behavioural analysis	73
4.3.2	EEG Analysis	77
4.4	Discussion	91
5	Hierarchical diffusion model analysis of perceptual decision making	95
5.1	Drift diffusion models	95

5.1.1	Drift diffusion models of decision making	98
5.1.2	Modelling behavioural data with drift diffusion models	100
5.1.3	Hierarchical diffusion models	103
5.1.4	Bayesian inference using Gibbs sampling	103
5.1.5	Advantages of diffusion models	104
5.2	Diffusion model analyses of the effect of learning on perceptual decision making	105
5.2.1	Model selection	105
5.2.2	Analysis of best fitting model	108
5.3	Diffusion model analyses of the effect of prior information on perceptual decision making.	111
5.3.1	Model selection	112
5.3.2	Analysis of best fitting model	113
5.4	Discussion	116
5.4.1	Investigating the possible effects of outlier	118
6	General conclusions	122
6.1	The neural locus of perceptual learning	123
6.2	Reinforcement learning models of perceptual learning	125
6.3	The role of prior expectation in perceptual learning	126
6.4	Drift diffusion models of perceptual learning	126
6.5	Limitations, extensions, & future directions	127
6.6	Conclusion & main contribution to the field of perceptual learning and decision making	128
	Appendices	130
A	Pilot experiments from Chapter 2	131
A.1	Overview	131
A.2	Materials & methods	132
A.2.1	Stimuli	132
A.2.2	General behavioural paradigm	132

A.2.3	Pilot experiment set 1	132
A.2.4	Pilot experiment set 2	133
A.2.5	Main pilot experiment	133
A.3	Results	136
A.4	Discussion of pilot studies	138
B	Pilot experiments from Chapter 4	139
B.1	Overview	139
B.2	General experimental design	140
B.3	Pilot experiment 1	141
B.3.1	Reaction time analysis	145
B.3.2	Discussion of Pilot Experiment 1	148
B.4	Pilot experiment 2	149
B.4.1	Accuracy analysis	150
B.4.2	Reaction time analysis	153
B.4.3	Conclusion from Pilot 2	156
B.5	Pilot experiment 3	157
B.5.1	Accuracy analysis	157
B.5.2	Reaction time analysis	159
B.5.3	Conclusion from Pilot experiment 3	160
B.6	General conclusions from pilot experiments	160
	References	162

List of Figures

1.1	An example of decision making in a psychology experiment.	2
1.2	Images illustrating the role of context on perception.	17
2.1	Examples of images used in the experiment.	26
2.2	Experimental design and behavioural analyses	37
2.3	Post-sensory effects of perceptual learning.	41
2.4	Enhanced readout of post-sensory decision evidence.	44
2.5	Probability of face response as a function of the decision variable.	48
3.1	Kahnt's reinforcement learning model of perceptual learning	57
3.2	Reinforcement learning model for perceptual choices.	60
3.3	Reinforcement learning model for perceptual choices for each subject	61
3.4	Extended RL model analysis	64
3.5	Electrophysiological correlates of prediction error (PE).	66
4.1	Experimental design and behavioural analysis.	72
4.2	Effects of prior information on perceptual learning and decision making	78
4.3	Discrimination variable varies with pre-stimulus cue in perceptual decision making	86
4.4	Post-cue & pre-stimulus change in neural activity during perceptual decision making	89
5.1	Illustration of a random walk process.	97
5.2	Illustration of drift diffusion model.	99
5.3	Graphical model of the best fitting model.	108
5.4	Illustration of diffusion model fit.	110

5.5	Graphical model of best fitting model.	114
5.6	Illustration of model fit.	115
5.7	Plots of all reaction times to identify outliers.	119
A.1	Behavioural task in pilot experiments.	133
A.2	Stimuli task in pilot experiments.	134
A.3	Behavioural task in the main pilot experiment.	135
A.4	Accuracy and reaction time in pilot experiments.	136
B.1	Behavioural task in pilot experiments.	140
B.2	Pilot experiment 1: Response accuracy.	141
B.3	Pilot experiment 1: Reaction time	145
B.4	Pilot experiment 1: Reaction time, for both correct and incorrect responses.	146
B.5	Pilot experiment 2: Response accuracy.	149
B.6	Pilot experiment 2: Response accuracy per subject.	153
B.7	Pilot experiment 2: Reaction time.	154
B.8	Pilot experiment 3: Response accuracy.	157
B.9	Pilot experiment 3: Reaction time, shown for accurate and inaccurate responses	159

Acknowledgements

Foremost, I would like to thank my supervisor, Dr Marios G. Philiastides, for all the guidance, encouragement and advice he has provided me throughout these four years. I will always be grateful for being part of his lab and for all the learning opportunities he gave me during this time, and for allowing me to become gradually more independent with my research. It has been a honour to be your PhD student at the Institute of Neuroscience & Psychology.

Also, I want to thank my second supervisor Prof. Lars Maku, and my PhD progress reviewers, Dr Guillaume Rousselet and Prof. Christoph Kayser, for always giving me important feedback that helped me to keep learning.

I also want to show my gratitude to Dr Yuka Sasaki and Prof. Takeo Watanabe, both of Brown University, for the discussions we had concerning my research, and especially for their *News & Views* in the journal *Nature Human Behaviour* that was about my first publication (i.e., Diaz, Queirazza, & Philiastides, 2017).

Thanks to all the PhD students and postdocs from the Institute of Neuroscience & Psychology and the others who were participants in my experimental studies. I want to thank them for their patience for coming in to the lab on three consecutive days, and never missing a day or coming late. Thanks especially to Essi for helping me when I needed it on the long days of EEG data collection.

I want to thank all the incredible people I met during my PhD who made my time at Glasgow extraordinary, especially all our lab members (Sabina, Elsa, Essi, Gabby, Andrea, Emma, Filippo, Guss, Leon and Tim) and to all my great roommates in Room 635 (Kasia, Kevin, Alex and Fei).

I want thank Martin Lomas, who has been a honorable role model, professionally & personally, when I worked with him as a social worker in Madrid. Without his encouragement, I wouldn't have taken the path of doing research in Psychology.

Also, I want to deeply thank the following people in Universidad Autónoma de Madrid: Dr. Jose Manuel Martinez, Dr. Jose Manuel Igoa, Dr. Carmen Almendros, Prof. Jose

Manuel Dols, Dr. Pablo Adarraga, Dr. Manuel Suero, Dr. Ricardo Olmos from the faculty of Psychology, and especially to Prof.Raul Villar, from the faculty of Physics. In different ways, they all have been important key influences in my research career, and all the learning opportunities I was offered during my psychology years have helped my knowledge to keep growing. It has been an honour being your student.

Finally, to all my loved ones, especially to Mark. I am most fortunate to have you all in my life.

List of Publications

- **Diaz, J. A.**, Queirazza, F., & Philiastides, M. G. (2017). Perceptual learning alters post-sensory processing in human decision making. *Nature Human Behaviour*, 1, Article Number: 0035 (2017). doi:10.1038/s41562-016-0035
- Philiastides, M. G., **Diaz, J. A.**, & Gherman, S. (2016). Spatiotemporal characteristics and modulators of perceptual decision-making in the human brain. In J.C. Dreher & L. Tremblay (Eds.), *Decision neuroscience - Handbook of Reward and Decision Making*. Academic Press.

The Diaz et al. (2017) paper above was subject the following editorial commentary in *Nature Human Behaviour*:

Sasaki, Y. & Watanabe, T. (2017). When perceptual learning occurs. *Nature Human Behaviour*, 1, Article number: 0048 (2017). doi:10.1038/s41562-017-0048

Author's Declaration

I declare that, except where explicit reference is made to the contribution of others, that this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or at any other institution.

List of Acronyms

- A1** Primary Auditory Cortex. 13
- ACC** Anterior Cingulate Cortex. 7, 15, 16, 124
- AIC** Akaike Information Criterion. 73–75, 84–87, 92, 107, 143, 147, 150–152, 154, 155, 157–160
- BIC** Bayesian Information Criterion. 47, 61, 62, 74, 86, 87
- CPP** Central Parietal Positivity. 8, 9, 128
- CRT** Cathode Ray Tube. 149
- DDM** Drift Diffusion model. 5–7, 9, 18, 19, 21, 24, 93, 95–105, 109, 111, 122, 127, 128
- DIC** Deviance Information Criterion. 74
- EEG** Electroencephalography. 8–10, 18, 19, 23, 24, 27–29, 32, 49, 50, 59, 63–65, 67, 69, 70, 72, 73, 77, 83, 84, 87–89, 91–94, 105, 118, 122, 123, 125–128, 138
- ERP** Event Related Potential. 8, 9, 128
- FFA** Fusiform Face Area. 18
- FMRI** Functional Magnetic Resonance Imaging. 6–8, 10, 12, 15, 18, 50, 51, 56
- HDDM** Hierarchical Drift Diffusion model. 21, 103, 105, 106, 111, 112, 116, 127, 128
- HPD** High Posterior Density. 48, 49, 109, 110, 114–116
- JAGS** Just Another Gibbs Sampler. 103, 104, 107
- LBA** Linear Ballistic Accumulator. 6, 7
- LIP** Lateral Interparietal. 5, 6, 9, 15, 16, 18, 54, 65, 68, 124
- MCMC** Markov Chain Monte Carlo. 104, 107, 112

MEG Magnetoencephalography. 8, 12, 18

MT Medial Temporal area. 2, 15, 18, 54, 65

NHPS Non-Human Primates. 16, 18, 24, 50, 53, 68, 117, 124

PET Positron-Emission Tomography. 12

PPA Parahippocampal Place Area. 18

RDK Random Dot Kinematogram. 2, 54

RL Reinforcement Learning. 15, 20, 53, 54, 57–62, 65–67, 125

ROC Receiver Operating Characteristic. 29, 31, 123

SC Superior Colliculus. 18

v1 Primary Visual Cortex. 12–14

v2 Secondary Visual Cortex. 13

v4 Visual Cortex Area 4. 14

WAIC Watanbe Akaike Information Criterion. 107, 108, 112, 113

Chapter 1

General introduction

In psychology and neuroscience, decision making is, generally speaking, the commitment to some categorization or classification of data, often followed by an action. While this very general description could apply to anything from the choice of a career or profession, to the choice of who to vote for in an election, for the most part, the term *decision making* in the context of psychology and neuroscience entails a relatively fast choice, i.e. on the order of hundreds to thousands of milliseconds, from a relatively constrained set of possibilities, and often on the basis of relatively constrained observable data. Examples of decision making of this kind include deciding if the following string of letters

uhnerts

is a word in the English language (see Ratcliff et al., 2004, for an example of a decision theoretic treatment of the lexical decision task), or after viewing the image shown in Figure 1.1 deciding if it contains a bottle of red wine (see Strayer & Kramer, 1994, for a decision theoretic description of visual memory search). In these examples, the choices are constrained to be binary, the response times will be on the order of hundreds of milliseconds, and the relevant data is clearly delimited.

A special case of the above type of decision making, and the one that is the focus of this thesis, is *perceptual decision making*. A perceptual decision is the classification of sensory, or sensory-motor, information usually followed by an overt behavioural response. Examples are easy to come by: A digital chirp is recognised (i.e., classified) as notification on our phone and we then pick up our phone to read the message; when passing someone



Figure 1.1: An example of decision making in a psychology experiment. This is an example image that could be used in a visual memory search experiment. The task could be, for example, to remember whether image contained a bottled of red wine.

on the street, we recognize that they are our colleague and we exchange greetings; when driving towards an intersection, we recognize the street sign as a Yield sign and slow down and cautiously drive through the intersection. In all these mundane examples, we experience auditory, visual and other sensory information, classify that information as belonging of some category, and then respond accordingly¹.

Perceptual decision making has received scientific attention only relatively recently. One of the seminal papers on this topic, if not *the* seminal paper, was by Newsome, Britten, and Movshon (1989). In that paper, Newsome et al. (1989) present the close correspondence between psychophysical measures of monkeys' decision making in the classification of the direction of motion of a Random Dot Kinematogram (RDK) and the response profiles of individual neurons in the Medial Temporal area (MT) area. Indeed,

¹Decision making and perceptual decision making are related to categorization. In both cases, stimuli are classified into categorically distinct classes. However, categorization, as the term is usually used in cognitive science, is more than recognition. It usually refers to the formation of semantic taxonomies, and how stimuli or items are assigned to categories within this taxonomy. For example, recognizing a hammer as a hammer would be an example of visual recognition, which is decision making according to our definition here. On the other hand, recognizing a hammer as an instance of a tool, and a tool as an instance of artefact, is more than visual recognition, and concerns how an object is assigned to a semantic categories or semantic hierarchies. Ultimately, we view recognition or decision making as a vital activity or component of categorization, but categorization goes beyond mere recognition.

they claim that the response profiles of a small number of individual neurons may be more sensitive and accurate in the measurement of visual motion than the monkey itself. This implies that perceptual decisions may be determined largely by the responses of a small number of cortical neurons.

1.0.1 Statistical decision theory

All decision making can be described from a statistical point of view. According to this perspective, we have the observed or available data \mathcal{D} , the set of possible classifications or choices $c_1, c_2 \dots c_K$, and a reward or loss function \mathcal{L} that gives the payoff or loss for any given choice. The probability that data \mathcal{D} is classified by choice c_k can be described in terms of Bayesian inference:

$$P(c_k|\mathcal{D}) = \frac{P(\mathcal{D}|c_k)P(c_k)}{\sum_{k=1}^K P(\mathcal{D}|c_k)P(c_k)}.$$

Here, $P(\mathcal{D}|c_k)$ is the probability of observing data \mathcal{D} given that true class is c_k , and $P(c_k)$ is the baseline probability of the class being c_k . Put another way, the *relative* probability of choice c_k versus choice c_l is given by

$$\frac{P(c_k|\mathcal{D})}{P(c_l|\mathcal{D})} = \frac{P(\mathcal{D}|c_k)}{P(\mathcal{D}|c_l)} \times \frac{P(c_k)}{P(c_l)}.$$

This shows that the relative probability of one choice versus another is the product of the relative evidence in favour of one choice over another and the relative prior or baseline plausibility of one choice over another.

While $P(c_k|\mathcal{D})$ gives the (subjective) probability that c_k is the correct choice, whether the subject in fact ultimately chooses c_k must take into account the payoff for choosing c_k . We can use \mathcal{L}_{kl} to denote the payoff (which could be a negative value) for choosing c_k when the correct choice is c_l . The *expected* loss for making choice c_k is

$$\sum_{j=1}^K \mathcal{L}_{kj}P(c_j|\mathcal{D}),$$

where we calculate the payoff for making choice c_k under all possible values of the correct choice, and then calculate the average payoff by weighting the payoff by the probability that each possible classification is correct. Therefore, it is optimal to choose c_k over c_l if

$$\sum_{j=1}^K \mathcal{L}_{kj}P(c_j|\mathcal{D}) > \sum_{j=1}^K \mathcal{L}_{lj}P(c_j|\mathcal{D}).$$

1.0.2 Sequential sampling models

We can consider the process of making a decision by a human or an animal as a process of evaluating evidence, and then when the evidence is above a certain threshold making a response. This is most easily appreciated in the binary choice case. In this case, we have two choices c_1 and c_2 . As mentioned above, the evidence in favour of c_1 relative to c_2 is

$$v = \frac{P(\mathcal{D}|c_1)}{P(\mathcal{D}|c_2)}.$$

Here, we can treat v as a subjective *decision variable*. If it is greater than 1, then there is more evidence for c_1 . If it is less than 1, then there is more evidence for c_2 .

As a consequence of the following relationship

$$\frac{P(c_1|\mathcal{D})}{P(c_2|\mathcal{D})} = \frac{P(\mathcal{D}|c_1)}{P(\mathcal{D}|c_2)} \times \frac{P(c_1)}{P(c_2)},$$

if $v > \frac{P(c_2)}{P(c_1)}$, then

$$P(c_1|\mathcal{D}) > P(c_2|\mathcal{D}).$$

As such, if the decision variable v crosses the threshold $\frac{P(c_2)}{P(c_1)}$, then it is optimal for the decision maker to choose c_1 as their response, assuming a symmetric payoff matrix (i.e. $\mathcal{L}_{11} = \mathcal{L}_{22}$ and $\mathcal{L}_{12} = \mathcal{L}_{21}$). More generally, if

$$v > \frac{(\mathcal{L}_{22} + \mathcal{L}_{12})P(c_2)}{(\mathcal{L}_{11} + \mathcal{L}_{21})P(c_1)},$$

then c_1 is the optimal choice, and if

$$v < \frac{(\mathcal{L}_{11} + \mathcal{L}_{21})P(c_1)}{(\mathcal{L}_{22} + \mathcal{L}_{12})P(c_2)},$$

then c_2 is the optimal choice.

From this, we can see that optimal decision making can be simply and accurately modelled by representing the relative evidence in favour of one decision over another by single scalar variable v , and then committing to one choice when v exceeds one upper threshold, and committing the alternative choice when v exceeds a lower threshold.

The previous description assumes that the evidence in favour of c_1 over c_2 has one fixed value. Of course, we can extend this idea to the case of where the evidence is building up over time. This is clearly a realistic and natural extension. For example,

when we observe an image, we do not instantaneously acquire all the information in the image, but rather we experience this information gradually as the eye moves over the image and our visual system processes this information (by way of its feed forward and feed back neural connections). According to this description, we now have a variable v_t that represents the evidence in favour of c_1 over c_2 on the basis of the information available so far at time t . However, whenever v_t crosses an upper or lower threshold, a commitment to one decision or the other is made. An illustration of an example of this process is provided on page 99 in Chapter 5 (Note that we devote all of Chapter 5 to the consideration of models of this nature.).

Mathematical models of decision making whereby a scalar random variable changes its value on the basis of experienced data until an upper or lower threshold criterion is met are known as sequential sampling models. Well established sequential sampling models include the Drift Diffusion model (DDM) (Ratcliff, 1978; Ratcliff & McKoon, 2008), which we describe in more detail in Chapter 5, the Linear Ballistic Accumulator (Brown & Heathcote, 2008), the leaky accumulator model (Usher & McClelland, 2001), and many more. In all these models — while differences exist between them that affect their interpretation and their degree of fit to the data — the general principles are the same: Evidence accumulates gradually until a threshold is met, and this signifies the commitment to a decision.

Sequential sampling models and cognitive neuroscience

Sequential sampling models have proved exceptionally good at modelling the speed and accuracy of the behavioural responses in decision making (see, e.g., Ratcliff & Smith, 2004; Ratcliff et al., 2004; Ratcliff & McKoon, 2008; Ratcliff & Van Dongen, 2011, for detailed reviews). Moreover, sequential sampling models have also provided particularly compelling accounts of the neuroscience of decision making. For example, Roitman and Shadlen (2002) have shown that Lateral Interparietal (LIP) area neurons, which have been shown to initiate saccades in motion discrimination tasks, build up their firing rates over time in accordance with the signal-to-noise ratio in the available data, and the saccadic response is made when a fixed threshold of firing rate activity is reached. In other words,

these decision related LIP neurons show the accumulation to threshold characteristics of all sequential sampling models.

In human neuroscience, sequential sampling models are routinely used to model the influence of external and internal variables and how they modulate decision making. For example, in sequential sampling models, accuracy of response and speed of response are inextricably coupled with one another and can not be treated as independent response variables. According to these models, arbitrarily high accuracy can always be achieved by increasing response times; arbitrarily fast responses can always be made when accuracy is sacrificed. Likewise, internal or external demands for greater accuracy must come at the expense of speed, and vice versa. Initial investigations of the neural basis of the speed versus accuracy modulation of decision making include Forstmann et al. (2008); Ivanoff, Branning, and Marois (2008). In Forstmann et al. (2008), in an Functional Magnetic Resonance Imaging (fMRI) based study, subjects were required to indicate whether a cloud of moving dots was moving on average to the left or right. In one condition, subjects were cued to increase speed, while in another they were cued to increase their accuracy. Analysis of behavioural data using a Linear Ballistic Accumulator (LBA) model (a two barrier sequential sampling model that is functionally similar to the DDM) shows that these experimental manipulations were almost perfectly accounted for by the changing of the distance between response barriers. fMRI analysis showed that the cueing for increased speed is associated with activity in the striatum and the pre-supplementary motor area, brain areas that are known to be involved in voluntary motor planning (Shima & Tanji, 1998). Moreover, individual differences in the level of activation of the striatum and pre-supplementary motor area was associated with individual variation in the estimated barrier change according to the LBA model. In Forstmann et al. (2010), again using a moving dot task where instructions varied between an emphasis on speed or on accuracy, showed that the strength of connection between the presupplementary motor area and striatum predicted the efficiency with which subjects could change their response threshold (as measured by an LBA model analysis). In related work, van Maanen et al. (2011) used a trial by trial variation in speed versus accuracy tradeoff. On trials where a speeded response was required, a positive relationship between trial by trial variation in the barrier height of the LBA model was found to correlate with activity in presupplementary motor

area and dorsal Anterior Cingulate Cortex (ACC). On trials where increased accuracy was demanded, the trial by trial variation in the LBA barrier was correlated with activity in ACC proper. Collectively, these results strongly imply that the speed and accuracy trade-off modulation is well modelled by changing the barrier height in accumulation to bound sampling models and furthermore that the neural basis of this system in the frontostriatal network.

Also motivated by the theoretical perspective of sequential sampling models, a study by Basten, Biele, Heekeren, and Fiebach (2010) has investigated how the costs and benefits of a particular decision can be modelled as an evidence accumulation process. Specifically, Basten et al. (2010) used an fMRI study to test the hypothesis that, when making perceptual decisions, evidence in favour of the potential costs of a decision is accumulated in parallel with evidence in favour of the decision's potential benefits and that the decision is eventually made when the difference between these two evidence accumulators passes a criterion threshold. The experiment task involved the presentation of colored shapes where subjects could accept or reject these stimuli in order to maximize their total reward. Each color and each shape had an associated monetary reward or monetary cost. During training, the colors and the shapes were presented separately and subjects learned the monetary values associated with each color and each shape. As such, during the experiment task itself, subjects must integrate their knowledge of the potential benefits of a particular color-shape combination against its potential costs. Reaction time and accuracy to accept or reject an offer both increased in accordance with the total payoff or total cost, respectively, of a decision. fMRI analysis identified a potential system underlying perceptual decision making as involving the accumulation of cost-benefit difference. For example, the ventral striatum and amygdala were identified as the sources of the benefits and costs, respectively, associated with the stimuli. The ventro-medial and left dorsolateral prefrontal cortex were shown to be involved in the integrating of the information being calculated by the striatum and amygdala in order to calculate a difference signal. Finally, the bilateral middle intraparietal sulcus was shown to be involved in the accumulation of this difference signal in order to reach the decision boundary.

While Basten et al. (2010) have demonstrated how the role of reward on decision making may be understood in terms of a DDM and have identified a potential neurophys-

iological basis for this system, it is still unclear whether the effects of changes in reward and costs primarily affect the earlier sensory processing stages or at later post-sensory processing stages. To address this question, an Electroencephalography (EEG) study by Blank, Biele, Heekeren, and Philiastides (2013) investigated the temporal dynamics of the role of punishment on decision making. Subjects in this experiment performed a face-car visual discrimination task, where the stimuli occurred at two levels of visual noise. Separate blocks in the experiment were associated with distinct punishment values (i.e., zero, medium, high) for incorrect responses. Behavioural analysis revealed that accuracy, but not reaction time, increased with punishment level. Single trial EEG analysis was used to identify EEG temporal component that best distinguish between the zero punishment and high punishment conditions. Unlike temporal components that distinguished between the high and low visual noise conditions, the punishment related temporal components appeared later in the trial, specifically around 400ms after stimulus onset or around 100ms prior to response, indicating that punishment affects later or post-sensory processing. In addition, the magnitude of the punishment temporal components correlate with the extent of the behavioural effects of the punishment conditions and trial by trial changes in the alpha and gamma EEG bands were predictive of the magnitude of the punishment components. Collectively, these results were taken to indicate that punishment affects attention and motivation which leads to more efficient accumulation of evidence to reach a decision.

1.0.3 EEG and the neural signatures of decision making

While Magnetoencephalography (MEG) and EEG have poor spatial resolution, their temporal resolution is on the order of milliseconds. As such, they have the potential to measure the time-course and temporal characteristics of human decision making to an extent not possible by, for example, fMRI.

Using a novel experimental paradigm, O'Connell, Dockree, and Kelly (2012) identify the neural signature of a *accumulation-to-bound* decision making variable in the human brain. In particular, the Central Parietal Positivity (CPP) Event Related Potential (ERP) activity was shown by O'Connell et al. (2012) to index the temporal integration of sensory evidence for decision making (see also Kelly & O'Connell, 2013). Similar to the

behaviour of LIP neurons described above (see, Roitman & Shadlen, 2002), the CPP ERP increases with signal quality and upon reaching a common threshold, a behavioural response is made.

An alternative approach to measuring the temporal correlates of human decision making has been taken by Philiastides et al (see, for example, Philiastides & Sajda, 2006; Philiastides et al., 2006a, 2006a; Ratcliff et al., 2009b). This approach aims to find a basis vector onto which the EEG signal is projected that will maximally discriminate between the EEG patterns in trials of one stimulus type compared to those of the other. On a two alternative forced choice visual discrimination task, specifically a face versus house image discrimination task, Philiastides et al identified two time points best discriminated between the stimulus categories. These are referred to as *temporal* EEG components. First, there is the *early* component that occurred around 170 ms after stimulus presentation and second there is a *late* component that occurred around 300 ms post-stimulus. The late component is identified as the *decision variable*. In comparison to the early component, the late component occurs later in time with increased task difficulty (Philiastides et al., 2006a). When the images are coloured red and green and the task is to discriminate red from green, the late component is reduced in magnitude to almost zero, while the early one is not affected. Also, in comparison to the early component, the late component is a significantly better predictor of trial-by-trial changes in the rate of evidence accumulation (i.e. drift rate) in a DDM (Ratcliff et al., 2009b). Taken together these findings indicate that the early component is a face recognition component that encodes the incoming sensory evidence, whereas the late component is a decision variable indicating the build up of decision-relevant evidence. (In Section 2.1, pages 23-24, we return to this work on EEG temporal components by Philiastides and Sajda (2006); Philiastides et al. (2006a, 2006a); Ratcliff et al. (2009b) and provide additional details and descriptions.)

1.0.4 Interim summary & introduction to the specific topics considered in this thesis

In the foregoing, we have described how human and animal decision making, and particularly perceptual decision making, can be described in terms of statistical decision theory,

which leads naturally to sequential sampling models of decision making. These models have been particularly effective models of the behavioural responses and neuroscience of decision making. In human neuroscience, while perceptual decision making has been studied using fMRI, and sequential sampling models have been applied to the description of fMRI decision making studies, EEG signals are at present more effective measures of the accumulation-to-bound characteristics of human perceptual decision making. In particular, the late temporal component identified using the single-trial EEG discriminant analysis pioneered by Philiastides et al has been shown to be a *neural signature* of perceptual decision making.

Accordingly, in this thesis, our aim is to address some open questions with respect to human perceptual decision making using the theoretical framework of sequential sampling models and the experimental paradigm of measuring temporal components in single-trial EEG discriminant analysis. The two open questions that we will address in this thesis are 1) how does learning or training affect perceptual decision making, and 2) how do prior experiences interact with learning to affect perceptual decision making.

1.1 The effects of learning on perceptual decision making

Perceptual learning, as defined by Goldstone (1998), involves any long-lasting changes to an organism's perceptual system that improve its ability to respond to its environment. As described by Law and Gold (2008a), it involves long-lasting improvements in our ability to detect, discriminate or identify sensory stimuli. Put less formally, perceptual learning involves any improvement in the speed or accuracy of an organism's perceptual system that results from learning or direct experience from the environment. A simple example of perceptual learning occurs in a visual discrimination task where noisy images must be classified as belonging to one of two mutually exclusive classes, namely faces or cars. In a task such as this, the subject's ability to quickly and accurately discriminate between the two classes of objects improves with practice with the task.

As defined above, perceptual learning is inextricably related to perceptual decision making. The outcome of perception, according to the descriptions just given, is that

of detecting, discriminating, identifying or responding to sensory stimuli. These are all examples of perceptual decision making, and as such, perceptual learning necessarily involves (though it is not limited to) improvements in the speed and accuracy of perceptual decision making. It is obvious, however, that, decision making is a general cognitive phenomenon that is also found in tasks, such as economic decision making, that are not just perceptual in nature. In addition, perception itself involves more than just decision making — decision making is essentially the final outcome of perception. An open question, therefore, and one that is the focus of this section, is the extent to which perceptual learning is due to improvements in general decision making processes or due to improvement in sensory abilities that are (informationally and temporally) prior to perceptual decision making.

Despite the obvious prevalence of perceptual learning in natural cognition, the underlying neural plasticity is still not well understood. We can identify at least three major hypotheses about the nature and location of perceptual learning in the brain. Learning may occur *early* in the perceptual system, such as in primary visual cortex. Learning may occur at *intermediate* levels in the perceptual system, for example, higher areas of the visual cortex. Finally, learning may occur *late* in, or even beyond, the perceptual system, specifically involving changes in the perceptual decision making system. We will review the evidence for these three hypotheses in turn.

1.1.1 “Early” perceptual learning

Until recently, the majority of studies on perceptual learning pointed to the conclusion that perceptual learning is occurring early in the perceptual system. This body of evidence can be sub-divided into psychophysical, neuroimaging and electrophysiological work.

Many psychophysical studies of perceptual learning in the 1980s and 1990s have shown that perceptual learning is often highly specific to the location and other properties of the stimuli used (see, for example, Sagi & Tanne, 1994, for an overview). This lends support to the hypothesis that perceptual learning arises early in the perceptual system. For example, Fiorentini and Berardi (1980) demonstrated that perceptual learning is specific to both orientation and spatial frequency of the gratings used in a discrimination task. In a hyperacuity task, Poggio, Fahle, and Edelman (1992); Fahle and Edelman

(1993) have shown that discriminative learning is highly specific to the location or orientation of stimulus in the visual field. Karni and Sagi (1991) have shown that in a visual texture discrimination task, learning is specific to the retina location of the stimulus, its orientation and also that there is no inter-ocular transfer of learning. Crist, Kapadia, Westheimer, and Gilbert (1997) reported that learning in a visual discrimination task is also specific to the location, orientation and geometric arrangement of the stimuli. Ball and Sekuler (1987) have reported improvements in motion discrimination that is specific to the direction of motion. Although these results are clearly behavioural in nature, given that the receptive fields of neurons early in the visual cortex are more tuned to specific locations in the visual field, as well as to orientation and spatial frequency, these results imply that the perceptual learning in these tasks is occurring early in the visual system.

Corroborating neuroimaging data for the hypothesis that perceptual learning occurs early in the perceptual system is also available. For example, Furmanski, Schluppeck, and Engel (2004) used fMRI to measure neural signals before and after a month long perceptual training period where subjects learned to detect oriented patterns. They observed an increase in responsiveness in Primary Visual Cortex (V1) cortex that correlated with improvement in behavioural performance. Similar results were obtained by Schwartz, Maquet, and Frith (2002) who used fMRI to identify neural correlates of perceptual learning in a texture discrimination task. The textured stimuli were shown to quadrants of the visual field. Using fMRI to compare activity in the retinotopic areas corresponding to the same quadrants of trained and untrained eyes, they showed higher activity in those areas corresponding to the trained eyes. Comparable results were found in the auditory domain. For example, using MEG, Pantev et al. (1998); Elbert, Pantev, Wienbruch, Rockstroh, and Taub (1995) found enhanced activity in early somatosensory and auditory cortices in trained musicians when compared to non-musicians, and the extent of activity in the musician correlated with the duration of their experience. Using Positron-Emission Tomography (PET), Molchan, Sunderland, McIntosh, Herscovitch, and Schreurs (1994) found increased activity in the primary auditory cortex after an associative learning task involving auditory stimuli, and Seitz and Roland (1992) found that increased blood flow in the primary somatosensory area as a consequence of a finger movement sequence was learned.

Direct electrophysiological recording of sensory neurons also supports the hypothesis that perceptual learning occurs early in the perceptual system. In Recanzone, Merzenich, Jenkins, Grajski, and Dinse (1992), adult owl monkeys were trained to detect differences in the frequency of a stimulus applied to their hands. They compared the primary somatosensory hand areas (area 3b) of the trained hands to the same areas in untrained hands. They also compared this somatosensory area in monkeys whose hands were stimulated but who did not perform a tactile discrimination task. Using electrophysiological recordings, the cortical representations of the trained hands were substantially more complex in topographic detail than the representations of monkey's untrained hands or of the passively stimulated hands. In a related study, Recanzone, Schreiner, and Merzenich (1993) trained owl monkeys for several weeks to discriminate differences in the frequencies of sequentially presented tones. Again, using direct electrophysiological recordings, the tonotopic organization of Primary Auditory Cortex (A1) in trained monkeys was compared to that in monkeys who heard the same stimuli, but did not engage in the discrimination task. The sharpness of tuning, and the latency of the response were greater for the trained monkeys when compared to control monkeys.

As a whole, these results clearly indicate that perceptual learning involves changes in the representation or processing of information in early stages in the perceptual system. This may involve, for example, changes in the receptive fields, tuning properties and activity rates of neurons in primary visual, somatosensory or auditory cortices.

1.1.2 “Intermediate” perceptual learning

Despite the body of evidence, just reviewed, for perceptual learning involving changes in the response properties of neurons in early parts of the perceptual system, some studies have failed to identify changes of this nature in otherwise similar experiments. For example, using electrophysiological recordings, Ghose, Yang, and Maunsell (2002) showed that, after long term training in a visual discrimination task with monkeys, the receptive fields of V1 and Secondary Visual Cortex (V2) neurons in trained and untrained regions were indistinguishable. Likewise, using a bisection visual discrimination task, Crist, Li, and Gilbert (2001) showed that the location, size and orientation selectivity of receptive fields were indistinguishable in trained and untrained monkeys. Yang and Maunsell

(2004) trained macaque monkeys to perform an orientation discrimination task at specific retinal locations and specific orientations. They showed that after training, neurons in Visual Cortex Area 4 (V4) but *not* in V1 had stronger responses compared to those in the untrained hemifield. These results were taken as providing one of the first demonstrations that the neural plasticity underlying perceptual learning occurs at an intermediate level in the visual cortex. More recently, work by Kuai, Levi, and Kourtzi (2013) has shown that higher ventral areas, such as the Lateral Occipital area, are more involved in the learning of decision templates in visual classification tasks.

Other recent evidence has shown that the location and orientation specificity of perceptual learning, reviewed above, may also not be unequivocal. For example, Xiao et al. (2008) has shown that training to discriminate a visual feature at one location could transfer to a second location if training with both locations occurred either simultaneously or immediately in succession. Xiao et al. (2008) argue that perceptual learning is central in its location and involves both discriminating specific stimulus features and learning to deal with stimulus-nonspecific factors like local noise at the stimulus location. When perceptual training with other stimulus locations had occurred, then the discrimination ability could transfer. In a similar manner, Ahissar and Hochstein (1997); Hochstein and Ahissar (2002); Ahissar and Hochstein (2004) have argued that generalization across stimulus characteristics such as orientation and retinal position is possible, that it readily happens when a task is not difficult, and that easy learning tasks followed by more difficult ones allow areas higher in the cortical hierarchy to guide generalization in lower areas.

1.1.3 “Late” perceptual learning

All the studies reviewed thus far have placed the locus of perceptual learning as being exclusively at some stage within the perceptual system itself, and the debate has centered on whether this stage is lower, such as V1, or higher, such as V4. A more intriguing possibility, and one that has been raised recently, is that perceptual learning occurs at higher or more central areas of the brain, beyond the perceptual system per se.

A recent primate based electrophysiology experiments by Law and Gold (2008a) has provided compelling evidence that perceptual learning arises from changes in decision making centres of the brain. Law and Gold (2008a) trained macaque monkeys to deter-

mine the direction of visual motion while making electrophysiological recordings from their MT and LIP areas. Area MT is well known to be involved in the representation of visual motion. The LIP is known to be involved in the control of eye-movement, particularly the execution of saccades. During the perceptual learning, the neurons in the LIP but not MT showed increasing responsiveness to motion, and there was a correlation between the neural responsiveness and the performance on the task. Law and Gold (2008a) conclude that these results imply that perceptual learning does not change how sensory information is *represented* in the brain, but rather how sensory representations are *interpreted*, particularly by the higher areas in the brain involved in decision making.

Law and Gold (2009) show that the results reported in Law and Gold (2008a) can be explained in terms of Reinforcement Learning (RL) (see, e.g., Sutton & Barto, 1998). They modelled the connections between the sensory neurons in the MT area and the decision making centre of LIP, and used a RL signal to modify the connections between these two areas. This explanation was corroborated by the work of Kahnt, Grueschow, Speck, and Haynes (2011) who trained subjects in a visual discrimination task where feedback was given. Their behavioural results are well modelled by a RL model that changes how the sensory information is interpreted. Using fMRI they found that activity patterns in the ACC, rather than in sensory areas, correlated with this learning.

In related work, S. Li, Mayhew, and Kourtzi (2009) also demonstrated that the locus of perceptual learning may be beyond the perceptual system itself. S. Li et al. (2009) used a categorical perception task whereby subjects had to decide if a stimulus (sampled from a continuum) was radial or spiral, and used fMRI to identify the brain regions that provide the neural signatures of the learning in this task. They identify that these regions are in frontal and higher occipitotemporal regions rather than in regions involved in signal detection or response execution such as primary visual or motor areas.

1.1.4 Interim summary concerning effects of learning on perceptual decision making

Although perceptual learning is well established phenomenon, its neural basis is not clear. In particular, there are open-questions about the locus of perceptual learning in the brain.

Initial empirical evidence, whether from psychophysics, neuroimaging or from electrophysiology, pointed to perceptual learning being largely based on changes in the earliest part of the perceptual system, e.g. primary visual cortex. Later work has shown that some of these earlier findings were unsound, and that sometimes higher areas of the perceptual system, and *not* lower areas, are the loci to perceptual learning. Moreover, there is no compelling evidence that areas beyond the perceptual system, and in particular, decision making centres of the brain such as LIP and ACC, are involved in perceptual learning, and that perceptual learning and more general reward based learning in the brain may have a common basis.

The recent research implicating general reward based learning and decision making mechanisms in perceptual learning has deep implications for our understanding of both perceptual learning and the role of decision making in perception. However, given that the clearest evidence for this comes from Non-Human Primates (NHPS) electrophysiology (i.e., Law & Gold, 2008a), it is currently not known whether and to what extent this revised interpretation extends to the human brain. Moreover, what neural mechanisms underlie these post-sensory processing changes remains unclear. The research that we carry out in this thesis directly addresses these main questions.

1.2 The role of prior expectation on perceptual decision making

In natural environments, perceptual stimuli do not occur in isolation or independently of their temporal and spatial context. Consider the two images shown in Figure 1.2. While both images depict familiar vehicles, clearly the image on the left is completely unnatural. Any time we observe a vehicle in the real world, as opposed to the artificial setting of psychophysics experiment, we observe it in the context of other objects and events. Moreover, these objects and events can not only have familiar spatial arrangements to one another, e.g. tables and chairs are usually arranged next to one another, a roof is always on top of a building, etc., objects can have familiar temporal arrangements too. For example, the flash of lightning precedes the crash of thunder, or the sound of your front door opening is often followed by the sight or sound of a familiar person. As such, the

spatial and temporal context potentially provide rich sources of information to facilitate the recognition perceptual stimuli.



Figure 1.2: **Images illustrating the role of context on perception.** The image on the left is unnatural because the object appears independently of any context.

The effect of prior expectation on perceptual processing has in fact been repeatedly demonstrated behaviourally (see, e.g., Bar, 2004; Oliva & Torralba, 2007, for reviews). For example, seminal work by Palmer (1975) showed that objects were identified more accurately when they are in familiar scenes, e.g. a toaster is recognized faster when seen in a kitchen. Similarly, Davenport and Potter (2004) show that objects are identified more accurately when they are in meaningfully consistent backgrounds, e.g. a ballerina is recognized more quickly when she is on a stage rather than when she is depicted on a football field. Auckland, Cave, and Donnelly (2007) show that objects are recognized more accurately when they are simply surrounded by meaningfully related objects. Auckland et al. (2007) claim that recognition of an object is affected by context not only when it is embedded in a consistent or familiar scene, but also when it is simply near other meaningfully related objects. These results are also not just limited to visual spatial relationships. For example, Saffran, Aslin, and Newport (1996); Aslin, Saffran, and Newport (1998); Fiser and Aslin (2002) show that temporal statistical relationships between auditory or visual stimuli can be learned even by infants.

Thus far, however, the neurobiological mechanisms underlying how prior expectations affect perceptual decision making are not well understood. Just as in the case of how and when perceptual learning occurs, which we considered in the previous section, there is

a debate concerning how and when prior expectation affects perceptual decision making. For example, Basso and Wurtz (1998) showed that in monkeys an increased probability of an upcoming perceptual stimuli, indicated by a perceptual cue, affected the build up of activity of Superior Colliculus (SC) neurons. The SC is known to be responsible for the control of saccadic eye movements. As such, these results imply that prior expectations are responsible for the preparation of a motor response to the stimulus. Similar results were obtained by de Lange, Rahnev, Donner, and Lau (2013), who used MEG with human participants, and showed that prior expectations based on pre-stimuli cues affect the build up oscillatory activity in the motor cortex. Related results were described by Albright (2012), who showed how, in monkeys, prior expectation activates MT neurons even before the stimulus occurs. Likewise, by using fMRI with humans, Puri, Wojciulik, and Ranganath (2009) showed that prior to stimulus onset, baseline activity in Fusiform Face Area (FFA) and Parahippocampal Place Area (PPA) was affected by prior expectation.

The literature just mentioned are consistent with a view that prior expectations affect baseline activity in sensory and motor areas *prior* to the stimulus onset. On the other hand, in a study using NHPs by Hanks, Mazurek, Kiani, Hopp, and Shadlen (2011), it was shown that the rate of firing of LIP neurons — as described above, LIP is known to integrate sensory signals towards a decision threshold — increased with increasing match between the prior expectation and the actual observed stimuli. This implies that prior expectation may affect the accumulation of evidence to reach a decision threshold, and not just affect the bias or baseline activity of the neural decision variable. Results obtained by Cravo, Rohenkohl, Wyart, and Nobre (2013), based on EEG analysis of human participants, also point to the role of prior expectation in affecting the efficiency of information accumulation, specifically by modulating the signal-to-noise gain of visual processing. Behaviourally consistent results with human subjects have been reported by Dunovan, Tremel, and Wheeler (2014), who demonstrated using a DDM that the primary role of prior expectation is to bias the rate of evidence accumulation favouring the more probable stimulus category.

1.3 Aims and outline of the thesis

Decision making, and perceptual decision making in particular, can be understood both theoretically and neurologically as a process of an accumulation of evidence to some threshold, at which point a commitment to a choice is made. This process can be examined in human subjects by analysing EEG data during perceptual decision making and identifying temporal components that are the neural signatures of the accumulation-to-bound decision making (see Philiastides & Sajda, 2006; Philiastides et al., 2006a, 2006a; Ratcliff et al., 2009b). It may also be statistically modelled using, for example, DDM (see Ratcliff & Smith, 2004; Ratcliff et al., 2004; Ratcliff & McKoon, 2008; Ratcliff & Van Dongen, 2011). Taken together, these provide us with a theoretical and experimental framework for the study of the neuroscience of human decision making.

Using this framework, we will investigate two related topics in the neuroscience of human decision making: How does learning affect perceptual decision making, and how does prior expectation interact with learning to affect decision making.

- Chapter 2 will describe our studies of the role of learning on perceptual decision making. As mentioned above, there are competing hypotheses about the nature and location of perceptual learning in the brain. The traditional view has been the *early* hypothesis that perceptual learning affects the sensory areas of the brains, for example, by modifying the receptive fields and interconnectivity of cells in primary sensory areas (see, e.g., Karni & Sagi, 1991; Sagi & Tanne, 1994; Recanzone et al., 1992, 1993). However, more recently Law and Gold (2008b) have argued that that learning affects how sensory representations are interpreted by the higher areas in the brain involved in decision making.

In the experiments described in Chapter 2, we will test if the account proposed by Law and Gold (2008b) also applies to human perceptual learning. We will do so by carrying out EEG recordings while subjects engage in a two-alternative forced choice visual object classification task, which they perform each day over the course of three days. We then use *single-trial* EEG analysis to describe how the temporal dynamics of decision making change as a consequence of perceptual learning. In particular, we will test how the early and late temporal components described

in Philiastides and Sajda (2006); Philiastides et al. (2006a, 2006a); Ratcliff et al. (2009b) are affected by training. We hypothesize that if the account proposed by Law and Gold (2008a) holds in humans, the late but not the early component should increase in strength and occur earlier in time with each day of training. This result would be consistent with a faster and more efficient accumulation of evidence.

- Chapter 3 deals with a specific mechanistic account of how learning affects perceptual decision making. This work follows from the work of Law and Gold (2009); Kahnt et al. (2011) who have applied RL to the study of perceptual learning. They argue that perceptual learning can be driven by a reinforcement signal that generates a selective readout of the sensory neurons that are most informative for any given perceptual decision. In Chapter 3, we apply related RL models to our perceptual learning results from the main experiment described in Chapter 2
- Chapter 4 will describe our studies of the interaction of prior expectation and learning on decision making. The behavioural task in the experiments described there will follow an identical paradigm used in our experiments described in Chapter 2: Subjects will perform a perceptual discrimination task, i.e., face/car discrimination, over the course of three days, using two different levels of stimulus noise. However, in addition, at the start of each trial one of three different cues will be presented. Each of these cues has an associated probability of the upcoming stimulus being a face or a car. We predict that when these cues are fully learned, consistent with Hanks et al. (2011), the late temporal component, but not the early temporal component, will increase its magnitude and occur earlier in time with the increasing match between the prior expectation and the actual observed stimulus. Following Dunovan et al. (2014), we predict a decrease in reaction time and an increase in accuracy with the increasing match between prior expectation and the observed stimulus. We also predict that the effects just mentioned will emerge gradually over the course of the three days. In other words, they will be relatively weak on the first day, stronger by the second, and strongest by the third. Finally, we predict that there will be an interaction between the effects of cue probability and training day. In particular, the effect of the prior expectation will *decrease* with increased

training. This particular prediction derives from Bayesian principle that prior information should influence decision making more when sensory evidence is weaker (e.g. Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006). If these predictions are met, they will show that learning affects decision making both by increasing the efficiency of the integration of sensory information but also by increasing the strength of top down expectations, (see, Summerfield & de Lange, 2014, for a review).

- Chapter 5 provides a Hierarchical Drift Diffusion model (HDDM) of the data from the main experiments described in Chapter 2 and Chapter 4. As mentioned above, the DDM is a sequential sampling model of decision making (see Ratcliff & Smith, 2004; Ratcliff et al., 2004; Ratcliff & McKoon, 2008; Ratcliff & Van Dongen, 2011). In this Chapter, we will model the response speed and choice data from the main experiments described in Chapter 2 and Chapter 4 using a HDDM. With respect to the first experiment, our particular aim is to determine how learning affects perceptual learning by examining which parameters of the diffusion model vary with training. With respect to our Chapter 4 experiment, our particular aim is to determine how prior probability of the upcoming stimulus, as revealed by the pre-stimulus cue, affects perceptual decision making by examining which parameters of the diffusion model vary with experimental predictor variables, but particularly with the pre-stimulus cue.

Chapter 2

Perceptual learning alters post-sensory processing in human decision making¹

2.1 Introduction

Perceptual learning involves any long-lasting improvements in our ability to detect, discriminate or identify sensory stimuli (see, e.g., Goldstone, 1998; Law & Gold, 2008a). As described in more detail in Chapter 1, the underlying neural plasticity of perceptual learning is still not well established. The objective of the study described in this chapter is to use electrophysiological data to obtain evidence either in favour of or against certain hypotheses about the nature of neural plasticity in perceptual learning.

As mentioned in Chapter 1, there are at least three major competing hypotheses about the nature and location of perceptual learning in the brain. Learning may occur *early* in the perceptual system, such as in primary visual cortex. Alternatively, learning may occur at *intermediate* levels in the perceptual system, for example, higher areas of the visual cortex. Finally, learning may occur *late* in, or even beyond, the perceptual system,

¹This work has been published as Diaz, J. A., Queirazza, F., & Philiastides, M. G. (2017). Perceptual learning alters post-sensory processing in human decision making. *Nature Human Behaviour*, Issue 1, Article Number: 0035. Some additional analyses that is not covered in this thesis can be found in this paper.

specifically involving changes in the perceptual decision making system. Of these three hypotheses, the dominant view has been the *early* hypothesis that perceptual learning affects the sensory areas of the brains, for example, by modifying the receptive fields and interconnectivity of cells in primary sensory areas (see, e.g., Karni & Sagi, 1991; Sagi & Tanne, 1994; Recanzone et al., 1992, 1993). However, recent primate based electrophysiology experiments (e.g., Law & Gold, 2008b) have provided compelling evidence against this dominant view. These results strongly imply that perceptual learning does not change how sensory information is *represented* in the brain, but rather how sensory representations are *interpreted*, particularly by the higher areas in the brain involved in decision making.

In the study described in this chapter, we test if this particular theoretical account of perceptual learning also applies to human perceptual learning. We do so by carrying out EEG recordings while subjects engage in a two-alternative forced choice visual object classification task, which they perform each day over the course of three days. We then use *single-trial* EEG multivariate discriminant analysis (see Section 2.2.5 for technical details) to provide a comprehensive account of how the temporal dynamics of decision making change as a consequence of perceptual learning. In this, we closely follow a paradigm established in previous studies using single trial EEG analysis of two alternative forced choice visual classification tasks (e.g., Philiastides & Sajda, 2006; Philiastides et al., 2006a; Ratcliff et al., 2009b; Philiastides et al., 2006a). In these studies, the same (or similar) two-alternative forced choice task was used, and using the same EEG discriminant analysis that we use here, Philiastides et al identified two temporally distinct neuronal components that discriminated between the stimulus categories: an *early* component that occurred around 170 ms after stimulus presentation and a *late* component that occurred around 300 ms post-stimulus.

Philiastides and Sajda (2006); Philiastides et al. (2006a); Ratcliff et al. (2009b); Philiastides et al. (2006a), established the following body of results that provide a comprehensive account of the neural significance of the *early* and *late* temporal EEG components:

- In Philiastides and Sajda (2006), it was shown that compared to the early component, the late component was a better predictor of behaviour. Specifically, a *choice probability analysis* indicated that there was a strong correlation between the neural

responses of the late component, but not the early one, and the subjects' behavioural judgements.

- In Philiastides and Sajda (2006); Philiastides et al. (2006a), it was shown that the late component, but not the early one, systematically shifted later in time with perceived task difficulty. For example, as noise increased in the stimuli, the onset time of the late component, but not the early one, occurred progressively later in time.
- In Ratcliff et al. (2009b), the late component was shown to be a significantly better predictor than the early component of trial-by-trial changes in the rate of evidence accumulation (i.e. drift rate) in a DDM.
- In Philiastides et al. (2006a), while the magnitudes of the early component were unaffected when the face and car stimuli were colored red or green and the task was switched to colour discrimination, the magnitudes of the late component were reduced almost to zero.

Taken together these findings indicate that the early component encodes the incoming sensory evidence, whereas the late component indexes *post-sensory, decision-relevant* evidence. These previous findings are very important for present purposes as they establish a body of facts against which to evaluate whether perceptual learning influences earlier versus later stages of decision making.

In the study described in this chapter, we will test how activity associated with each of these early and late EEG components is affected by training. We hypothesize that if perceptual learning primarily alters post-sensory encoding of decision evidence, as seen in NHPs (Law & Gold, 2008a), discrimination performance for our late but not the early component should systematically increase across the three training sessions. We specifically hypothesize that single-trial amplitudes of the late decision-related components will increase as a function of learning and systematically move backward in time consistent with a faster and more efficient accumulation of evidence. Moreover, our ability to exploit single-trial variability in the EEG will offer a mechanistic characterization of these effects by establishing whether improvements in discrimination are a result of gain modulation (i.e. amplification of the differential response) of the component amplitudes, a reduction in the trial-to-trial variability (i.e. noise) of the component amplitudes, or both.

2.2 Materials and Methods

2.2.1 Subjects

14 Subjects (7 Women and 7 men, age range 23-28 years) participated in this study. All were right handed, reported normal vision and no history of neurological problems. Informed consent was obtained from all participants in accordance with the guidelines of the Institute of Neuroscience and Psychology from Glasgow University. The study was approved by the College of Science and Engineering Ethics Committee at the University of Glasgow (CSE01353).

2.2.2 Stimuli

The stimuli were selected from the stimulus set described in e.g., Philiastides and Sajda (2006, 2007). The stimulus set was generated as follows. A set of 20 images of faces and 20 images of cars were selected from the Face Database of the Max Planck Institute of Biological Cybernetics² (Troje & Bühlhoff, 1996; Blanz & Vetter, 1999). Each image was 512×512 pixels, with 8 bits per pixel, and there were equal numbers of frontal and side (up to 45°) views. All images were equated for spatial frequency, luminance, and contrast. They all had identical magnitude spectra and their corresponding phase spectra were manipulated using the weighted mean phase (Dakin, Hess, Ledgeway, & Achtman, 2002) technique to generate a set of images characterized by their percentage phase coherence. For each image, a set of 13 noisy variants were created. The noise levels were described in terms of *coherence* and ranged uniformly from 20% (lowest coherence, highest noise) to 50% (highest coherence, lowest noise level). As such, in this stimulus set there were a total of $2 \times 20 \times 13 = 520$ images. We present examples of the stimuli used in Figure 2.1.

We selected two levels of sensory evidence for this study (32.5% and 37.5% phase coherence) that are known to yield performance spanning the psychophysical threshold, based on previous studies (Philiastides, Ratcliff, & Sajda, 2006b; Philiastides & Sajda, 2006).

²Available at <http://faces.kyb.tuebingen.mpg.de>.

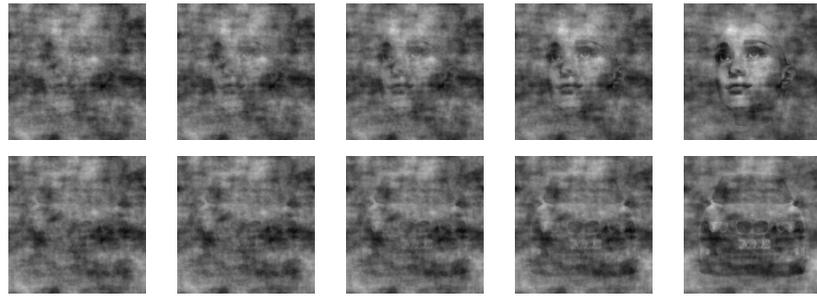


Figure 2.1: Examples of a face image (upper row) and a car image (lower row), each at 5 coherence levels: From left to right, their coherence values are 30%, 35%, 40%, 45% and 50%.

A Dell Precision Workstation (Intel Core 2 Quad) running Windows 7 (64 bit) with an ATI FirePro 2270 graphics card and PsychoPy 1.80 presentation software (Peirce, 2007) controlled the stimulus display. Images were presented on a Dell 2001FP TFT monitor (resolution, 1600 x 1200 pixels; refresh rate, 60 Hz).

2.2.3 Behavioural Task

Subjects were presented with the stimuli just described and performed a two alternative forced choice classification task whereby they classified each image as either a face or a car. Subjects sat a distance of 75 cm from the computer monitor and each image subtended approximately 6x6 degrees of visual angle. On each trial, a blank screen was displayed for a random duration that ranged uniformly between 1.00 to 1.50 seconds. The stimulus image was then presented for 50 milliseconds and subjects were given up to 1250 milliseconds to make their classification response, which was done using a USB button box using their right hand's index (for face response) and middle (for car response) fingers. Subjects received visual feedback for each response. For a correct response, a tick was presented. For an incorrect response or if the subject did not respond within the 1250 interval, a cross was presented. The feedback duration lasted 500 milliseconds. The trials were presented in 4 blocks of 72 trials, with a 60 second rest period between each block. The entire experiment lasted approximately 20 minutes. In Figure 2.2a, we provide the task diagram of the behavioural task.

Each subject performed this task on three consecutive days. In other words, there were 42 experimental testing sessions in total, i.e. 14 subjects each tested on three separate

days, with the experiment taking place at the same time on each day so that there was 24 hours between each session for all subjects. On the first day, subjects performed a practice session of the face/car classification task but with a different set of face and car images.

2.2.4 EEG Data Acquisition

Subjects performed the task on three consecutive days, in a dark and soundproof room. As they performed the task, their EEG was recorded with a 64 channel Ag/AgCl scalp electrode actiCAP EEG system (Brain Products GmbH, Gilching, Germany). The active ground electrode was placed just below Pz electrode of the International 10-20 system method. The active reference electrode was placed on the left mastoid. The impedance was always below 5kOhm for each subject on each day.

The EEG signal was acquired at 1000hz at an analogue band pass of 0.02-250hz. The EEG data was re-referenced to the average of all channels. Prior to beginning the experiment on each day, subjects performed an eye-movement calibration task. They were instructed to blink naturally at a white cross at the centre on the screen of the computer, and after a few seconds they made a right-left, up-down, movement following the white cross. The timing of these visual cues was recorded with the EEG. These eye blinks and movement artefacts were removed by using a principal component analysis based approach described in Parra, Spence, Gerson, and Sajda (2005). Finally, we baseline corrected the EEG data, with the baseline interval defined as the 100 ms prior to stimulus onset.

2.2.5 Single Trial EEG Data Analysis

We performed single trial discriminant analysis in order to describe how the temporal dynamics of decision making changes as a consequence of perceptual learning. Here, we closely followed the paradigm established in previous studies by Philiastides et al (see, for example, Philiastides & Sajda, 2006; Philiastides et al., 2006a, 2006a; Ratcliff et al., 2009b) that was used to identify the neural signatures of perceptual decisions. This method provides a non-invasive means to study the temporal evolution of perceptual decision making. In this approach, a single-trial analysis of the EEG signal is performed as

follows. The EEG activity at time t on a single trial, i.e. trial i is denoted by \vec{x}_{it} . Note that \vec{x}_{it} is a column vector with as many elements, denoted by K , as there are EEG electrodes. In our case, $K = 64$. Corresponding to \vec{x}_{it} , we have z_i , which is a binary variable that indicates some cognitive or perceptual characteristic of trial i . For example, in a two alternative forced choice visual discrimination task, z_i could indicate the identity of the visual stimulus shown on trial i . Our aim is to now find a *basis vector* \vec{w}_t that best discriminates the EEG vectors on those trials for which $z_i = 0$ from the EEG vectors on the trials for which $z_i = 1$. Each \vec{x}_{it} can be projected onto \vec{w}_t to give use a scalar variable

$$y_{it} = \vec{w}_t' \vec{x}_{it}.$$

Note that y_{it} is a summary representation of the activity in \vec{x}_{it} that best discriminates one class (denoted by $z_i = 0$) from the other class (denoted by $z_i = 1$). The value of \vec{w}_t can be found by logistic regression. In particular, by finding \vec{w}_t that maximizes

$$L_t = \prod_{i=1}^N p_{it}^{z_i} (1 - p_{it})^{1-z_i},$$

where

$$p_{it} = \frac{1}{1 + \exp(-y_{it})},$$

and

$$y_{it} = \vec{w}_t' \vec{x}_{it}.$$

From what we have just described, \vec{w}_t is a *weighting vector* of the EEG activity \vec{x}_{it} . In other words, y_{it} is a weighted sum of the EEG activity \vec{x}_{it} with the weighting coefficients given by \vec{w}_t . We may also obtain a neuroanatomical interpretation of \vec{w}_t by treating y_{it} as a source and viewing \vec{x}_{it} as derived from the value of y_{it} as follows:

$$\vec{x}_{it} = \vec{a}_t y_{it} + \vec{\epsilon}_{it}.$$

Here, \vec{a}_t is a *coupling vector* that multiplies the discriminating component y_{it} to give us the (non-noise) contribution of the discriminating component to the sensory activity \vec{x}_{it} . As described in, e.g., Parra et al. (2002, 2003), the value of \vec{a}_t can be calculated by a least-squares solution as follows:

$$\vec{a}_t = \frac{\mathbf{X}_t \mathbf{y}_t}{\mathbf{y}' \mathbf{y}},$$

where \mathbf{X}_t is the $K \times N$ matrix formed by concatenating the column vectors \vec{x}_{it} , for all $i \in 1 \dots N$ trials, and \mathbf{y}_t is the $N \times 1$ vector formed by concatenating y_{it} for all $i \in 1 \dots N$ trials. This \vec{a}_t is known as the *sensor projection* (e.g., Parra et al., 2003) or *scalp projection* (e.g., Parra et al., 2002; Philiastides & Sajda, 2006) that can be visualized as *scalp maps* that show the neuroanatomical significance of the discriminating component.

Accordingly, we used a logistic regression based multivariate discriminant analysis to identify the EEG patterns that reliably distinguish between the two alternative stimulus types in the forced choice task. This analysis was applied separately to each subject on each day as follows. For any trial i for an arbitrary subject on an arbitrary day, \vec{x}_{it} denotes the $K = 64$ dimensional EEG signal at timepoint $t \in \{-100, -90 \dots -10, 0, 10 \dots 1000\}$ after the stimulus onset in that trial. Note that the timepoints are in steps of 10ms apart. The set of $\delta + 1$ signals

$$[\vec{x}_{it-\delta/2} \cdots \vec{x}_{it+\delta/2}],$$

with $\delta = 50$, were treated as $\delta + 1$ independent observations of the EEG signal at post stimulus timepoint t . For each \vec{x}_{it} , we define a classifier label as

$$z_{it} = \begin{cases} 1 & \text{if the stimulus of trial } i \text{ is a face,} \\ 0 & \text{if the stimulus of trial } i \text{ is a car.} \end{cases}$$

We can quantify the performance of the discriminant analysis classifiers by the area under the Receiver Operating Characteristic (ROC) curve, which we label as A_z . Thus, for each subject on each day and at each time point t , we have an A_z value. The significance of this A_z statistic was computed using bootstrapping whereby we randomly permute classification labels.

In more detail, the ROC curve allows us to quantify the performance of any binary classifier. Recall that, as noted on page 28, y_{it} is a summary representation of the EEG activity that best discriminates one of our two classes from the other. This y_{it} deterministically maps to a probability — specifically, the probability that the class to which the EEG activity corresponds is the one coded as class 1 — by the equation

$$p_{it} = \frac{1}{1 + \exp(-y_{it})}.$$

Thus, in general with logistic regression, for every pattern of stimulus activity, we have a probability that this activity represents or is indicative of class 1. With this probability

value, we can then choose any threshold $0 \leq \Phi \leq 1$ to create a decision rule to classify any stimulus as belonging to, or corresponding to, class 0 or 1 as follows:

$$\text{classification decision for stimulus } it = \begin{cases} 0, & \text{if } p_{it} \leq \Phi, \\ 1, & \text{if } p_{it} > \Phi \end{cases}$$

For the classifier defined by any particular value of Φ , there will be a *false positive* rate and a *true positive* rate. The false positive rate is the number of false positives (stimuli classified as 1 when their true class is 0) divided by the number of true negative cases (the number of stimuli that truly belong to class 0). The true positive rate is the number of true positives (stimuli correctly classified as 1) divided by the number of true positive cases (the number of stimuli that truly belong to class 1). As an example, consider the fictitious data presented in the following table:

	stimulus	truth	p	$\Phi = 0.10$	$\Phi = 0.30$	$\Phi = 0.50$	$\Phi = 0.70$	$\Phi = 0.90$
1	-0.88	0	0.87	1	1	1	1	0
2	-0.75	1	0.85	1	1	1	1	0
3	1.38	0	0.29	1	0	0	0	0
4	0.24	1	0.63	1	1	1	0	0
5	0.11	1	0.67	1	1	1	0	0
6	1.20	0	0.34	1	1	0	0	0
7	-0.46	1	0.80	1	1	1	1	0
8	0.64	1	0.51	1	1	1	0	0
9	0.42	1	0.58	1	1	1	0	0
10	0.78	0	0.47	1	1	0	0	0
False positive rate				1.00	0.75	0.25	0.25	0.00
True positive rate				1.00	1.00	1.00	0.33	0.00

In this example, our stimulus is one-dimensional data and each stimulus belongs to class 0 or 1, which is shown in the column labelled *truth*. We then fit a binary logistic regression to this data, and then obtain a predicted probability, for each stimulus, that it belongs to class 1. This is given by p in the table. Now, we create a classification rules as just defined using the thresholds $\Phi = 0.10$, $\Phi = 0.30$, $\Phi = 0.50$, $\Phi = 0.70$, $\Phi = 0.90$. For

each threshold, there is obviously a false positive rate and a false negative rate. We can see, for example, that when $\Phi = 0.10$, all cases are classified as belonging to class 1, leading to a true positive rate of 1.00 (i.e., all truly positive cases are correctly classified as belonging to class 1), but also a false positive rate of 1.00 (i.e., all truly negative cases are incorrectly classified as belong to class 1). On the other extreme, when $\Phi = 0.90$, all cases are classified as belonging to class 0, leading to a true positive rate of 0.00 (i.e., all truly positive cases are incorrectly classified as belonging to class 0), but also a false positive rate of 0.00 (i.e., no truly negative cases are incorrectly classified as belonging to class 1).

The ROC space is simply a 2 dimensional space, specifically a unit square, i.e., $(0, 1) \times (0, 1)$, where the x-axis represents the false positive rate and the y-axis represents the true positive rate. Any classifier defined a threshold in the manner just described, can be represented by a point in this space. For example, when $\Phi = 0.70$, the false positive rate is 0.25 and the true positive rate is 0.33, and so this classifier can be represented in the ROC space by the point $(0.25, 0.33)$. The ROC *curve* is then simply the curve through the ROC space defined by the set of points in ROC space corresponding to the classifiers defined by the continuum of thresholds from $\Phi = 0.00$ to $\Phi = 1.00$. This curve will necessarily start at the point $(0.00, 0.00)$ and terminate at $(1.00, 1.00)$. A curve close to identity line, i.e., where the y value is equal to the x value, indicates the logistic regression, on average, discriminates approximately at chance. On the other hand, an ROC curve extends close to the upper left value of the ROC space indicates that the logistic regression, on average, discriminates well between the two classes. Thus, we can quantify the performance of the logistic regression based discrimination analysis by calculating the area under the ROC curve, and we'll denote the value of this area by A_z . If A_z is close to 0.50, it indicates that the discrimination analysis is poor. On the other hand, if A_z is close to 1.00 it indicates that the discrimination analysis is highly accurate.

In order to assess the significance of the A_z value for any logistic regression, we can apply a permutation test, also known as a *bootstrapping* test (Parra et al., 2005). In this method, we randomly permute the class labels that are applied to the stimuli, and then recalculate the logistic regression discrimination analysis and its A_z value. This produces a set of sampled A_z that we will denote by $\tilde{A}_z^1, \tilde{A}_z^2 \dots \tilde{A}_z^B$, where B is the number of per-

mutations we apply. To obtain a $p = 0.05$ significance level for our A_z , we find the 95th percentile value of $\tilde{A}_z^1, \tilde{A}_z^2 \dots \tilde{A}_z^B$, i.e. the value above which lie 5% of the samples. In other words, if the null hypothesis were true, and the class labels are essentially random, the probability of observing an A_z value *as or more extreme* than the 95th percentile value of our permutation samples is 0.05. As such, any observed A_z that is greater than this threshold can be regarded as significant at the 0.05 level.

2.2.6 Multilevel Regression Analyses

In the analysis of the behavioural and EEG data, we use *multilevel*, or *random effects*, general and generalized linear models (Gelman & Hill, 2007). Given that we use multilevel models extensively below, we will now provide a general introduction to them.

In order to better appreciate the details of the multilevel models that we use, it is useful to consider the following simplified example. Assume that in an experiment, each of J subjects perform a speeded visual classification task, e.g., a face/car discrimination task, a total of n times. In total, therefore, across all subjects, we will have $N = J \times n$ experimental trials. The reaction times across these trials can be labelled $y_1, y_2 \dots y_N$; the true identity of the visual stimulus on each trial, i.e. whether it was a face or car, can be represented by the binary variables $x_1, x_2 \dots x_N$, with each $x_i \in \{0, 1\}$; and the identity of the subject on each trial can be represented by $h_1, h_2 \dots h_N$, with each $h_i \in \{1 \dots J\}$.

If there was absolutely no inter-subject variability in the relationship between the stimulus identity x_i and reaction time y_i , and assuming normally distributed trial by trial variability in reaction time, we could model this data simply by

$$y_i \sim N(b_0 + b_1 x_i, \sigma^2), \quad i \in \{1 \dots N\},$$

which is equivalent to

$$y_i = b_0 + b_1 x_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad i \in \{1 \dots N\}.$$

We can then use, for example, maximum likelihood estimation to infer the values of the parameters b_0 , b_1 and the trial by trial variance parameter σ^2 .

On the other hand, if there is inter-subject variability in the relationship between x_i and y_i , a multilevel model of this relationship is

$$y_i \sim N(\phi_{0[h_i]} + \phi_{1[h_i]}x_i, \sigma^2), \quad i \in \{1 \dots N\},$$

$$\phi_{0[j]} \sim N(b_0, \tau_0^2), \quad \phi_{1[j]} \sim N(b_1, \tau_1^2), \quad j \in \{1 \dots J\}.$$

Here, we are explicitly modelling inter-subject variability in the linear relationship between x_i and y_i in terms of normally distributed variation around the linear coefficients b_0 and b_1 . The parameters describing these variations are τ_0^2 and τ_1^2 , respectively. Note that in the limiting cases of $\tau_0^2 \rightarrow 0$ and $\tau_1^2 \rightarrow 0$, meaning no inter-subject variability, this multilevel model is identical the previous non-multilevel model above.

Given that $\phi_{0[j]}$ may be rewritten as $\phi_{0[j]} = b_0 + \beta_{0[j]}$, where $\beta_{0[j]} \sim N(0, \tau_0^2)$, and $\phi_{1[j]} = b_1 + \beta_{1[j]}$, where $\beta_{1[j]} \sim N(0, \tau_1^2)$, the above multilevel model is equivalent to

$$y_i = \underbrace{b_0 + b_1x_i}_{\text{fixed effects}} + \underbrace{\beta_{0[h_i]} + \beta_{1[h_i]}x_i}_{\text{random effects}} + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2),$$

$$\beta_{0[j]} \sim N(0, \tau_0^2), \quad \beta_{1[j]} \sim N(0, \tau_1^2), \quad j \in \{1 \dots J\}.$$

Note that here, the random effects coefficients have a mean of zero. As such, this multilevel model effectively models the relationship between x_i and y_i by a global linear equation $b_0 + b_1x_i$ that is subject to zero-mean normally distributed random variation across subjects. The global or fixed effects equation $b_0 + b_1x_i$ therefore represents the average population relationship between x_i and y_i , and is often the particular focus of an analysis.

It should also be noted that model fitting in multilevel models involves, for example, maximum likelihood estimation of the following parameters: b_0 , b_1 , τ_0^2 , τ_1^2 , and σ^2 . As such, the estimated values of $\beta_{0[j]}$ and $\beta_{1[j]}$, for $j \in \{1 \dots J\}$, play a role analogous to the values of ε_i , for $i \in \{1 \dots N\}$ in that they describe the random variation observed on individual trials. They are generally not explicitly analyzed unless attention needs to be paid to an individual data point or individual subject.

The inclusion of the random effects terms is to properly account for the inter-subject variability around these average effects, and we do so primarily so that the statistical significance of the average effects is calculated correctly, with failure to properly account for

inter-subject variability leading to inflated Type I error rates (Aarts, Verhage, Veenliet, Dolan, & van der Sluis, 2014, for further discussion of the necessity of multilevel analysis in neuroscience). In this sense, multilevel general linear models can be viewed as playing a similar to *repeated-measures* Anova models. However, they allow us to overcome the known statistical inadequacies of repeated-measures Anova analyses, (Baayen, Davidson, & Bates, 2008, for further discussion).

The significance of a single variable, or set of variables, in the multilevel regression models may be tested using a log likelihood ratio test. For example, to test the significance of a size K' subset of all K predictor variables (i.e. $K' \subseteq K$), we compare the log-likelihood of the model with all K predictors against the log-likelihood of the model without the K' subset. If we denote the log-likelihood of the model with all K predictors by \mathcal{L}_1 and the log-likelihood of the model with the K' subset by \mathcal{L}_0 , then under the null hypothesis that all coefficients corresponding to the K' predictors are simultaneously zero,

$$-2 \times (\mathcal{L}_0 - \mathcal{L}_1) \sim \chi_{\text{df}=\Delta_\theta}^2$$

where Δ_θ is the difference in the number of parameters in the two models. In other words, under the null hypothesis, -2 times the difference of the log likelihoods of the two models is distributed according to a χ^2 distribution whose degrees of freedom is Δ_θ . The parameter counts in each model are based on both the number of predictors as well as the size of the covariance matrix (Bates, 2010). For example, consider the following three multilevel linear models:

- Example multilevel model I:

$$y_i = b_0 + \beta_{0[h_i]} + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad \beta_{0[j]} \sim N(0, \tau_0^2), \quad j \in \{1 \dots J\},$$

- Example multilevel model II:

$$y_i = b_0 + b_1 x_i + \beta_{0[h_i]} + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad \beta_{0[j]} \sim N(0, \tau_0^2), \quad j \in \{1 \dots J\}.$$

- Example multilevel model III:

$$y_i = b_0 + b_1 x_i + \beta_{0[h_i]} + \beta_{1[h_i]} x_i + \varepsilon_i, \\ \varepsilon_i \sim N(0, \sigma^2), \quad \beta_{0[j]} \sim N(0, \tau_0^2), \quad \beta_{1[j]} \sim N(0, \tau_1^2), \quad j \in \{1 \dots J\}.$$

Model I has three degrees of freedom, corresponding to the one fixed effect intercept term, the one random effects intercept term, and the residual variance term. Model II has four degrees of freedom, three corresponding to those terms also in Model I, and then an additional degree of freedom corresponding to fixed effect slope term. Model III has five degrees of freedom, four of which correspond to those terms also in Model II, and then an additional degree of freedom corresponding to the random effects slope term. It should also be noted that it is possible to model the correlation between random intercepts and random slope terms. In this case, there is one additional degree of freedom in the model corresponding to the covariance between the random intercept and random slope. In general, if we have h random effects variables, if we allow for non-zero correlation between each of these, the random effects collectively take $\frac{h \times (h+1)}{2}$ degrees of freedom.

Note that the likelihood ratio test just mentioned can always be used to test the significance of a categorical variable with L levels, which is coded using $L - 1$ binary *dummy* codes, or an interaction between two categorical variables with L and L' levels, which can be coded using $(L - 1) \times (L' - 1)$ binary dummy codes.

The likelihood ratio test affords us a general means to test hypotheses. As mentioned, it can be used to test hypotheses related to the role of individual predictor variables, including categorical predictor variables, or sets of predictor variables. It can also be used identically with multilevel and non-multilevel general linear or generalized linear (for example, binary logistic models) models. For this reason, it is our default means for hypothesis testing and model comparison. In some analysis, other methods of hypothesis tests are also possible. For example, in some cases, to assess whether a particular predictor variable has a reliable effect on an outcome variable, we may use the p-value for the null hypothesis test that the coefficient for the predictor is significantly different to zero. Equivalently, we may calculate whether the, for example, 95% confidence interval for the predictor contains the value of zero. In many cases, tests of this nature will lead to identical or near identical results to the corresponding likelihood ratio test. However, depending on circumstances (for example, the type of statistical model, the type of predictor variable, etc), and depending on the assumptions of the models and the assumptions of the tests, the results of these different hypothesis test and model comparison methods may not be identical. This is, in principle, inevitable given that there ultimately many different

ways of statistically testing hypotheses or comparing models, and no one of them is a definitive method.

We fit all multilevel linear and logistic regression models using the lme4 package

<http://cran.r-project.org/web/packages/lme4/index.html>

in the R statistical language

<http://www.r-project.org>.

2.3 Behavioural analysis results

In Figure 2.2b-c, we have plotted the average accuracy and average response times on each day of training as a function of the coherence level in the stimulus. On 1.5% of trials across all subjects and training days, subjects did not respond within the 1250ms response interval, and these were excluded from the analysis. For the trials for which choices were made, we calculated the average accuracy on each day and for each of the levels of coherence, averaging over the trials and all subjects. These are plotted in Figure 2.2b. For trials with correct responses, we calculated the average reaction times on each day and each of the two levels of coherence, again averaging over all trials and all subjects. These are shown in Figure 2.2c.

We modelled the accuracy data using a multilevel logistic regression analysis, and modelled the reaction times for accurate responses using multilevel linear regression. In these analyses, the predictor variables were the coherence of the stimulus (treated as a binary variable with values "low" and "high") and training day (treated as a continuous variable whose values range from 1 to 3), and the interaction of coherence level and training day. The identity of the subject was treated as a random effect varying the slope and intercept of the regression. More formally, for the accuracy analysis, we used the following model:

$$\log \left(\frac{P(z_i = 1)}{1 - P(z_i = 1)} \right) = \underbrace{b_0 + b_1 x_{1i} + b_2 x_{2i} + b_3 x_{1i} x_{2i}}_{\text{fixed effects}} + \underbrace{\beta_{0h_i} + \beta_{1h_i} x_{1i} + \beta_{2h_i} x_{2i} + \beta_{3h_i} x_{1i} x_{2i}}_{\text{random effects}}$$

$$\beta_j \sim N(0, \Sigma), \quad \text{for } j \in 1 \dots J,$$

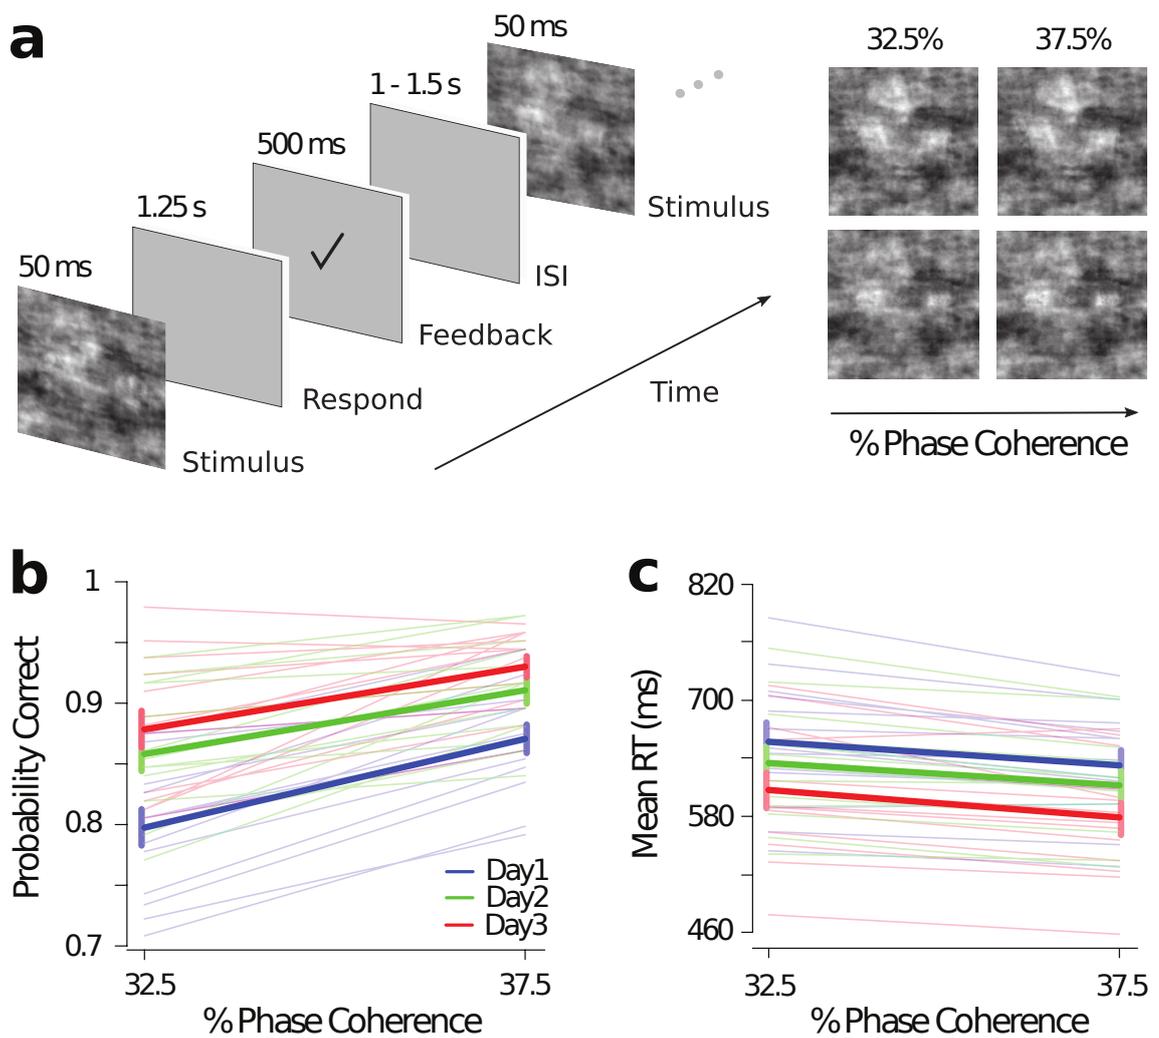


Figure 2.2: **Experimental design and behavioural analyses.** Schematic representation of the experimental paradigm (a). Subjects had to categorize a noisy image presented for 50 ms as a face or a car and indicate their choice with a button press within 1,250 ms following the stimulus presentation. Feedback was then presented for 500 ms (a tick or a cross, for a correct or an incorrect response, respectively) followed by an inter-stimulus interval (ISI) that varied randomly between 1 and 1.5s. Subjects performed this task on three consecutive training days. A sample face image (upper row) and car image (lower row) at the two levels of phase coherence used in the task (32.5% and 37.5%) are shown on the right. Shown in (b) are the proportion of correct choices, and shown in (c) are the mean reaction times (RTs), as a function of the three training days (Day 1: blue; Day 2: green; Day 3: red) and the two levels of phase coherence of the stimuli, averaged across subjects. Faint lines represent individual subject data. Error bars represent standard errors across subjects.

where z_i is a binary variable indicating the accuracy of the subject on trial i , and where x_{1i} and x_{2i} represent the value of the coherence variable and the value of the training day variable on trial i , respectively. For the reaction time analysis, we used the following model:

$$y_i = \underbrace{b_0 + b_1x_{1i} + b_2x_{2i} + b_3x_{1i}x_{2i}}_{\text{fixed effects}} + \underbrace{\beta_{0h_i} + \beta_{1h_i}x_{1i} + \beta_{2h_i}x_{2i} + \beta_{3h_i}x_{1i}x_{2i}}_{\text{random effects}} + \varepsilon_i,$$

$$\varepsilon_i \sim N(0, \sigma^2), \quad \text{for } i \in 1 \dots N \quad \beta_j \sim N(0, \Sigma), \quad \text{for } j \in 1 \dots J,$$

where y_i signifies the reaction time of the subject on trial i , and the predictor variables x_{1i} and x_{2i} are identical to those of the logistic regression model for accuracy. (Note that in both models, to avoid notational clutter, we denote the coefficients of the fixed and random effects using identical symbols, i.e. b_0, b_1, b_2, b_3 and $\beta_0, \beta_1, \beta_2, \beta_3$, respectively.).

In order to better appreciate the details of these two multilevel models, it useful to consider the arrangement of the data used in both models:

accuracy	reaction time	coherence	training day	subject identity
z_1	y_1	x_{11}	x_{21}	h_1
z_2	y_2	x_{12}	x_{22}	h_2
\vdots	\vdots	\vdots	\vdots	\vdots
z_i	y_i	x_{1i}	x_{2i}	h_i
\vdots	\vdots	\vdots	\vdots	\vdots
z_N	y_N	x_{1N}	x_{2N}	h_N

Here, each of J subjects perform a speeded visual classification task, e.g., a face/car discrimination task, a total of n times. In total, therefore, across all subjects, we will have $N = J \times n$ experimental trials. As such, z_i is a binary variable indicating the accuracy of the subject on trial i ; y_i is a continuous variable indicating the reaction time of the subject on trial i ; x_{1i} is a binary variable indicating the coherence value of the stimulus presented on on trial i ; x_{2i} is a continuous variable indicating the training day (i.e., 1, 2 or 3) on trial i ; finally, h_i is a categorical variable, i.e. $h_i \in \{1 \dots J\}$, indicating the identity of the subject on trial i .

Note that *subject identity*, signified by h_i , is *not* a predictor variable, but is a grouping variable, or random effect variable, in the multilevel model. In other words, we can understand the multilevel model as providing a distinct logistic or linear regression model for *each individual subject*, i.e. J distinct regression models. In each of these J regression models, we have *three* predictors: *coherence*, *training day* and the interaction thereof. However, these J regression models are coupled to one another due to the fact that the coefficients for each of these predictors are assumed to be sampled from a *global* distribution.

For each analysis, we used a log-likelihood ratio test, as described above, to test for the main effects and interaction effects. For accuracy, there were significant main effects for coherence level ($\chi^2_{df=1} = 28.08$, $p < 0.01$) and the day of the session ($\chi^2_{df=1} = 19.37$, $p < 0.01$). There was no interaction between coherence level and training day ($\chi^2_{df=1} = 0.16$, $p = 0.68$). For reaction time, there were significant main effects for coherence level ($\chi^2_{df=1} = 21.24$, $p < 0.01$) and the day of the session ($\chi^2_{df=1} = 8.92$, $p < 0.01$). There was no interaction between coherence level and training day ($\chi^2_{df=1} = 0.38$, $p = 0.54$). Note that for each one of these analyses, the degrees of freedom of the χ^2 statistic for the log-likelihood ratio test is $df = 1$. This is because we always compare the log-likelihood of the full model with a model with one less parameter. For example, to assess the interaction effect in the accuracy analysis, we compare the log-likelihood of the logistic regression model shown above to the log-likelihood of the following model:

$$\log \left(\frac{P(z_i = 1)}{1 - P(z_i = 1)} \right) = \underbrace{b_0 + b_1 x_{1i} + b_2 x_{2i}}_{\text{fixed effects}} + \underbrace{\beta_{0h_i} + \beta_{1h_i} x_{1i} + \beta_{2h_i} x_{2i} + \beta_{3h_i} x_{1i} x_{2i}}_{\text{random effects}}$$

$$\beta_j \sim N(0, \Sigma), \quad \text{for } j \in 1 \dots J.$$

As such, there is one less parameter, i.e. b_3 , in the smaller model³.

³As explained in Section 2.2.6, page 35, there may be alternative methods of testing some hypotheses of this nature. However, our default method for hypothesis testing will be the log-likelihood ratio test using the χ^2 statistic as this is the most generally and widely applicable method available.

2.4 EEG analysis results

The onset times of the early and the late components are obtained independently for each subject on each day. They are found by identifying the time-point corresponding to the *peak*⁴ A_z value in time windows from around 100ms-250ms (early component window) and 250ms-500ms (late component window), see Philiastides and Sajda (2006). In all cases, the late component was never within 100ms of the response time of the subject. For each subject on each day, having recorded the time of the early and late components, we then obtained the A_z values of these early and late component times.

In Figure 2.3a, we show an example trajectory of the y discriminating variable (introduced and described in Section 2.2.5, page 28) for the early (dashed lines) and late (solid lines) component windows for one example subject. In Figure 2.3b, we have plotted the A_z values of the discrimination analysis as a function of each time point in the trial on each day. In Figure 2.3c, we have plotted the A_z values of early and late components as a function of day, averaged over subjects. In Figure 2.3c, we also show the scalp topologies of the early and late components (calculated using the method described in Section 2.2.5, page 28). As can be seen, the late component's regions of highest activity are in more parietal and occipitoparietal areas, while the early component's regions of highest activity are more central. These results are in line with the topologies of the early and late components in, for example, Philiastides and Sajda (2006). The onset times of the early and late components are shown as a function of day for each subject in Figure 2.3d.

2.4.1 Analysis of peak A_z values

We analyzed the effect of the temporal component (treated as a binary variable with two levels, i.e., "early" and "late"), training day (as before, treated as a continuous variable with values in the range 1 to 3), and the interaction of temporal component and training

⁴Whenever we use the term *component onset time*, we generally are referring to the time-point of the *peak* of this component.

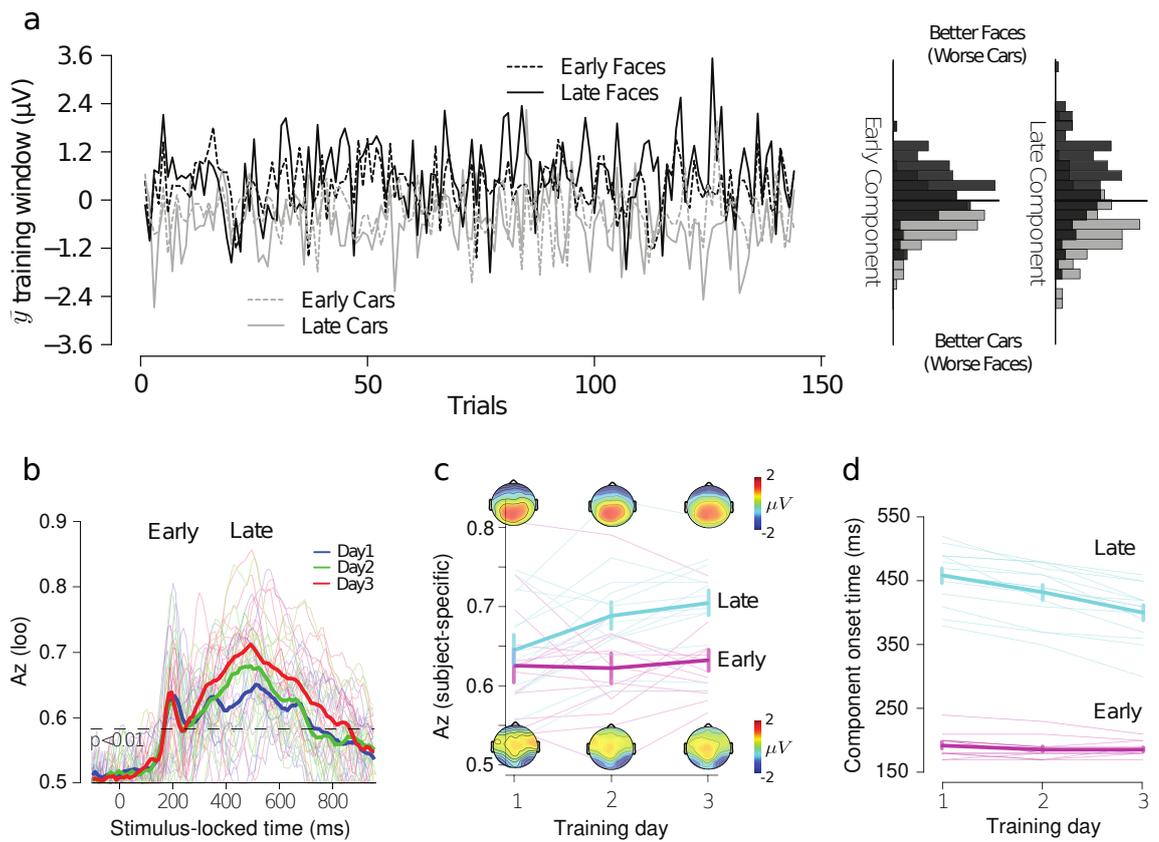


Figure 2.3: Post-sensory effects of perceptual learning. Single-trial discriminator amplitudes (y) for the early (dashed lines) and late (solid lines) component windows for faces (black) and cars (grey) at 37.5% phase coherence from a representative subject on the third training day (a). The component amplitudes are shown as histograms on the right, with a cutoff (the thick black line) to separate trials into positive versus negative amplitudes, indicating a higher likelihood of a face and a car trial, respectively. (b) Multivariate discrimination performance (A_z) during face-versus-car outcome discrimination of stimulus-locked EEG responses across the three training days (Day 1: blue; Day 2: green; Day 3: red), averaged across subjects, showing the presence of the early and late components. The dotted line represents the average A_z value leading to a significance level of $P = 0.01$, estimated by using a bootstrap test. Faint lines represent individual subject data. (c) Average discriminator performance and scalp topographies for the early (magenta) and late (cyan) components across the three training days estimated at the time of subject-specific maximum discrimination. Faint lines represent individual subject data. Error bars represent standard errors across subjects. (d) Average onset times for the early (magenta) and late (cyan) components across the three training days. Faint lines represent individual subject data. Error bars represent standard errors across subjects.

day, on the value of the A_z peak using the following multilevel linear regression model:

$$A_{z_i} = \underbrace{b_0 + b_1x_{1i} + b_2x_{2i} + b_3x_{1i}x_{2i}}_{\text{fixed effects}} + \underbrace{\beta_{0h_i} + \beta_{1h_i}x_{1i} + \beta_{2h_i}x_{2i} + \beta_{3h_i}x_{1i}x_{2i}}_{\text{random effects}} + \varepsilon_i,$$

$$\varepsilon_i \sim N(0, \sigma^2), \quad \text{for } i \in 1 \dots N \quad \beta_j \sim N(0, \Sigma), \quad \text{for } j \in 1 \dots J.$$

where $i \in \{1 \dots N\}$ indexes the data-points, $N = 14 \times 3 \times 2 = 84$ (i.e., two components on each day for each subject), A_{z_i} is the value of the peak of the A_z value, $x_{1i} \in \{1, 3\}$ is the training day, $x_{2i} \in \{0, 1\}$ is the early or late temporal component.

As before, we used a log-likelihood ratio test to test for the main effects and interaction effects. There was a significant main effect of component ($\chi^2_{df=1} = 5.04, p = 0.02$), a significant main effect of day ($\chi^2_{df=1} = 7.61, p = 0.01$), and a significant interaction between training day and component ($\chi^2_{df=1} = 7.46, p = 0.01$)⁵. Note that in the above model, the coefficients b_1 and b_2 are the coefficients for the main effect of training day and the main effect of temporal component, respectively, while the coefficient b_3 is the interaction effect for training day and temporal component. In particular, b_3 gives the *change* in slope of the effect of training day when we go from the early temporal component to the late temporal component. The estimated value of $b_3 = 0.03, t(40.99) = 2.77, p < 0.01$ ⁶. The significant value for b_3 here shows that the effect of training in the late component is significantly greater than that of the early component. In more detail, the maximum likelihood estimate of b_1 is 0.003, with its parametric bootstrap based 95% confidence interval being $(-0.01, 0.02)$ ⁷. On the other hand, the maximum likelihood estimate of b_3 is 0.03, as mentioned already, with its parametric bootstrap based 95%

⁵Note that as described in Section 2.2.6, and as was the case with the analyses described in Section 2.3, for each one of these analyses, the degrees of freedom of the χ^2 statistic for the log-likelihood ratio test is $df = 1$. This is because we always compare the log-likelihood of the full model with a model with one less parameter.

⁶The t-tests obtained from the multilevel model are for null-hypothesis tests on the coefficients b_1, b_2, b_3 , just as they would be in, for example, a non-multilevel general linear model. However, given the nature of these multilevel models, the *Welch-Satterthwaite* approximation, which is calculated using the R package <https://cran.r-project.org/web/packages/lmerTest/lmerTest.pdf>, is used to calculate the degrees of freedom. This approximation usually results in *non-integer* degrees of freedom.

⁷Note that the confidence interval for b_1 contains 0.00. This has happened despite the fact that the model with day as a predictor is a significantly better model, as measured by log-likelihood, than the model without day. This occurs simply because the effect of day on the outcome variable is close to zero for the

confidence being (0.01, 0.04). Thus, the maximum likelihood estimate of the slope of the effect of training day in the early component is 0.003, while for the case of the late component it is $0.003 + 0.03 \approx 0.03$.

Collectively, these results show that the A_z value for the late component, but not the early one, increased significantly in magnitude over the course of training.

2.4.2 Analysis of component onset times

Using an identical multilevel analysis to that used to analyze peak A_z values, we modelled the effect of temporal component, training day, and their interaction, on the onset times of the temporal components. Using a log-likelihood ratio test as before, there was a highly significant main effect of day, $\chi^2_{df=1} = 21.56$, $p < 0.01$; a highly significant main effect of component, $\chi^2_{df=1} = 51.30$, $p < 0.01$; a highly significant interaction between day and component, $\chi^2_{df=1} = 51.75$, $p < 0.01$. The maximum likelihood estimate of b_3 is -26.07 , $t(41.00) = -9.88$, $p < 0.01$. As in the case of the A_z peak analysis, this shows that the onset time of the late component occurs significantly earlier as an effect of training than does the early component.

The maximum likelihood estimate of b_1 is -3.21 , with its parametric bootstrap based 95% confidence interval being $(-9.00, 2.22)$. The maximum likelihood estimate of b_3 is -26.07 , with its parametric bootstrap based 95% confidence being $(-31.34, -20.44)$. Thus, the maximum likelihood estimate of the slope of the effect on onset time of the training day in the early component is -3.21 , while for the case of the late component it is $-3.21 - 26.07 = -29.28$.

2.4.3 Analysis of the discriminating component y

The improvement in subjects' visual discrimination that results from training over the three days, as evidenced by the behavioural data alone is consistent with multiple alternative underlying mechanisms. In particular, when viewed from the point of view of standard signal detection theory, improvements in discrimination as a result of learning could

early component, but non-zero for the late component. The effect of day on the outcome variable for the early component is given by b_1 , while the effect of day on the outcome variable for the late component is $b_1 + b_3$.

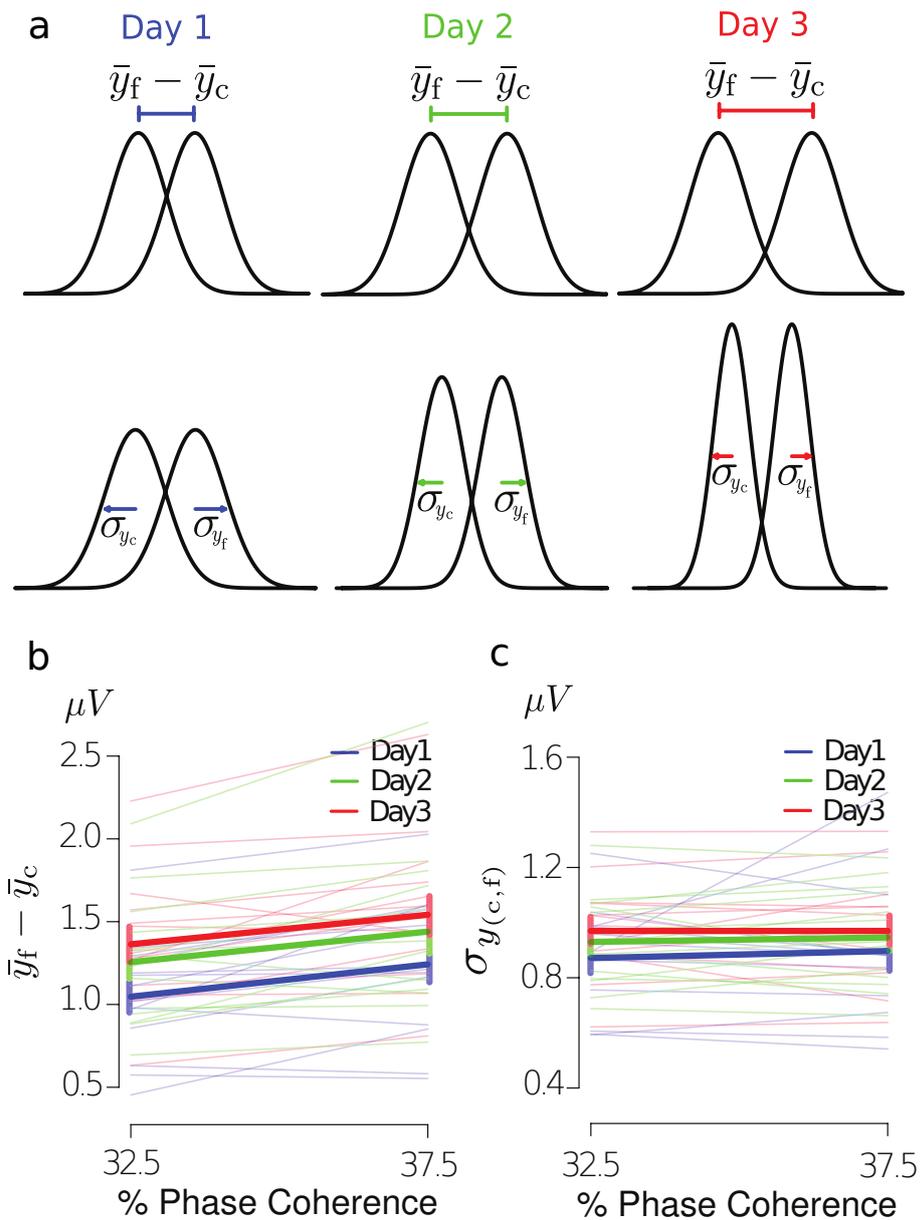


Figure 2.4: **Enhanced readout of post-sensory decision evidence.** (a) A schematic illustration of possible effects of learning on the distribution of single-trial discriminator amplitudes in the course of learning. Top: increases in the distance between the mean response for faces and cars. Bottom: reduction in the variance of the face and car responses. Both examples lead to a smaller overlap (more separation) between the face and car distributions. (b) Changes in the mean distance between the face and car distributions for the late component across the three training days (1: blue, 2: green, 3: red). (c) Changes in the variance of the face and car distributions for the late component across the three training days. The faint lines in b and c represent individual subject data, while the error bars represent standard errors across subjects.

be caused by either increasing the perceptual separability of the two stimulus classes, or alternatively by decreasing the noise and variability in the signals. We illustrate these two possibilities in the Figure 2.4a, respectively. Analyzing the y discriminating component, which effectively represents the neural decision variables, can potentially provide valuable insight into which of either of these two possibilities is underlying the effect of learning on perceptual decision making.

The average difference in the late discriminating component y for faces and cars as a function of coherence and days is shown in Figure 2.4b. Likewise, the average standard deviation of the late discriminating component as a function of coherence and days is shown in Figure 2.4c.

Separately for the early and late components, we used a multilevel linear model to analyze how the difference between the discriminating component for faces and for cars varied as a function of coherence and days, modelling inter-subject variability in these effects as random effects. For these two analyses, we used a multilevel linear model identical in form those used to analyze the magnitude and onset of the temporal components, i.e.

$$y_i = \underbrace{b_0 + b_1x_{1i} + b_2x_{2i} + b_3x_{1i}x_{2i}}_{\text{fixed effects}} + \underbrace{\beta_{0h_i} + \beta_{1h_i}x_{1i} + \beta_{2h_i}x_{2i} + \beta_{3h_i}x_{1i}x_{2i}}_{\text{random effects}} + \varepsilon_i,$$

$$\varepsilon_i \sim N(0, \sigma^2), \quad \text{for } i \in 1 \dots N \quad \beta_j \sim N(0, \Sigma), \quad \text{for } j \in 1 \dots J.$$

where, in this model, y_i signifies the difference in the values of the discriminating component, while $x_{1i} \in \{1, 3\}$ signifies the training day, and $x_{2i} \in \{0, 1\}$ signifies the level of coherence of the stimulus.

For the early component analysis, there was a significant effect for coherence $\chi^2_{df=1} = 16.10$, $p < 0.01$, but not for day $\chi^2_{df=1} = 0.11$, $p = 0.74$ nor for the interaction of day and coherence $\chi^2_{df=1} = 0.44$, $p = 0.51$. For the late component, on the other hand, there was a significant effect of day $\chi^2_{df=1} = 6.72$, $p < 0.01$, and coherence $\chi^2_{df=1} = 11.52$, $p < 0.01$, but no interaction between the two $\chi^2_{df=1} = 0.03$, $p = 0.86$.

An identical multilevel analysis was then used to analyze the effect of training and coherence on average standard deviation of y . For the early component, there was no significant effect of day $\chi^2_{df=1} = 0.24$, $p = 0.62$, coherence $\chi^2_{df=1} = 1.13$, $p = 0.29$, or interaction $\chi^2_{df=1} = 1.65$, $p = 0.20$. Likewise for the late component, there was no signifi-

cant effect of day $\chi_{df=1}^2 = 2.76$, $p = 0.10$, coherence $\chi_{df=1}^2 = 0.38$, $p = 0.53$, or interaction $\chi_{df=1}^2 = 0.25$, $p = 0.61$.

In sum, from these analyses, the main results are that the average distance between the discriminating components for faces and cars increases with learning. We do not, however, observe any significant change in the average standard deviation with learning.

2.4.4 Prediction of behavioural responses by early versus late component

We analyzed how well the values of the early and the late components predict the subjects' behavioural responses on each trial. For each trial from each subject, we will use $r_i \in \{0, 1\}$ to indicate if the *response* of the subject is “face” ($r_i = 1$), or “car” ($r_i = 0$). Likewise, on each trial, we have y_i^{early} and y_i^{late} , which are the values of y decision variable from the discriminant analysis at the early component and late components, respectively.

We used multilevel analyses to see if y_i^{early} or y_i^{late} is a better predictor of behavioural responses. Two models are used. In one, only y_i^{early} is used to predict behavioural responses. In the other, only y_i^{late} is used to predict behavioural responses. Then, the model fit statistics of the two models are compared. All data, across all subjects and all days, are used. We index the individual trials by i , i.e. i ranges from 1 to $J \times 3 \times 288$, where 288 is the number of trials per subject per day and J is the number of subjects. On trial i , $h_i \in \{1 \dots J\}$ indicates the identity of the subject. A multilevel logistic regression of how the probability of a face response varies as a function of y_i^{early} is

$$\log \left(\frac{\text{P}(r_i = 1)}{1 - \text{P}(r_i = 1)} \right) = \alpha_{h_i} + \beta_{h_i} y_i^{\text{early}}$$

where

$$\alpha_j \sim \text{N}(a, \tau_a^2), \quad \beta_j \sim \text{N}(b, \tau_b^2)$$

are random effects on the intercept and slope that account for inter-subject variability in the model. The corresponding multilevel logistic regression of how the probability of a face response varies as a function of y_i^{late} is

$$\log \left(\frac{\text{P}(r_i = 1)}{1 - \text{P}(r_i = 1)} \right) = \alpha_{h_i} + \beta_{h_i} y_i^{\text{late}}.$$

Fitting these two models using R, and calculating the Bayesian Information Criterion (BIC), we obtain

- BIC for early model: 14046.44.
- BIC for late model: 13087.51.

The BIC is defined as follows:

$$-2LL + \log(n)K,$$

where K is the number parameters in the model. Note that a *lower* BIC value indicates *better* model fit. In particular, according to Kass and Raftery (1995), we can calculate the approximate log of the Bayes Factor comparing the early model to the late model by

$$\log(\text{Bayes Factor}) \approx -(\text{BIC}_{\text{late}} - \text{BIC}_{\text{early}})/2.$$

A widely used rule (see, for example Kass & Raftery, 1995) is that if $2 \times \log(\text{Bayes Factor}) > 10$, this is *very strong evidence* that the late model is a better model of response choice than the early one. As such, our model comparison result is

$$-(\text{BIC}_{\text{late}} - \text{BIC}_{\text{early}})/2 = 479.47$$

and so $2 \times \log(\text{Bayes Factor}) = 958.93$, which is clearly much greater than 10 and so is very strong evidence that the late component is a predictor of the subjects' responses than is the early component.

In a complementary analysis, using the R package `MCMCglmm` that implements a Bayesian inference in multilevel model, we fit the model

$$\log\left(\frac{\text{P}(r_i = 1)}{1 - \text{P}(r_i = 1)}\right) = \underbrace{a + b_{h_i}^{\text{early}} z_i^{\text{early}} + b_{h_i}^{\text{late}} z_i^{\text{late}}}_{\text{fixed effects}} + \underbrace{\alpha_{h_i} + \beta_{h_i}^{\text{early}} z_i^{\text{early}} + \beta_{h_i}^{\text{late}} z_i^{\text{late}}}_{\text{random effects}},$$

$$\alpha_j \sim \text{N}(0, \tau_a^2), \quad \beta_j^{\text{early}} \sim \text{N}(0, \tau_{b_e}^2) \quad \beta_j^{\text{late}} \sim \text{N}(0, \tau_{b_l}^2)$$

where z_i^{early} and z_i^{late} are standardization of the y_i^{early} and y_i^{late} , defined as follows:

$$z_i^{\text{early}} = \frac{y_i^{\text{early}} - \langle y_i^{\text{early}} \rangle}{\text{sd}(y_i^{\text{early}})} \quad z_i^{\text{late}} = \frac{y_i^{\text{late}} - \langle y_i^{\text{late}} \rangle}{\text{sd}(y_i^{\text{late}})}$$

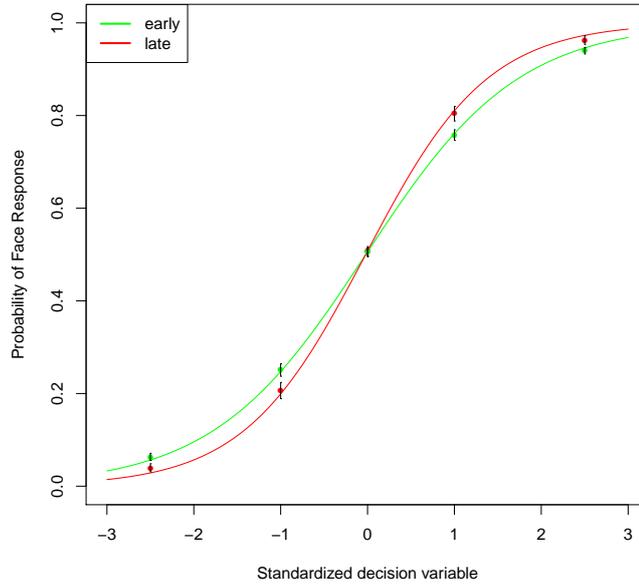


Figure 2.5: **Probability of face response as a function of the decision variable.** For each of the 14 subjects, the probability of a face response was modelled as a function of either the (standardized) early or late decision variable, using logistic regression. The points shown are the average, across subjects, probabilities of a face response at values of the decision variable $-2.50, -1.00, 0.00, 1.00, 2.50$. Also shown are the standard errors for these predictions. The curves are the average curves obtained from logistic regression analyses across all subjects.

Using these standardized variables, we can compare the compare β^{early} and β^{late} as both predictors are on a similar scale⁸. We provide a plot of the probability of a face response as a function of the standardized decision variable in Figure 2.5.

The Bayesian analysis allows us to easily calculate interval estimates for the coefficients and for differences in the coefficients. These interval estimates, known as High Posterior Density (HPD) intervals, can be interpreted as roughly analogous to confidence intervals. The estimates and 95% HPD intervals for b^{early} and b^{late} are

	Coefficient	95% HPD Interval
b^{early}	0.88	(0.75, 1.01)
b^{late}	1.24	(1.07, 1.45)

⁸In general, two coefficients in a regression model are not comparable unless the predictor variables to which they correspond have the same scale.

The average difference $b^{\text{late}} - b^{\text{early}}$ is 0.37 and the 95% HPD interval on this difference is (0.15, 0.59), indicating that the standardized value of the late component is a reliably better predictor of the probability of a “face” or “car” response than is the standardized value of the early component.

2.5 Discussion

In this work, we offer the first evidence from time-resolved electrophysiological signals in humans linking perceptual learning with post-sensory processing during a perceptual categorization task. Specifically, we showed that improvements in behavioural performance were accompanied primarily by late enhancements in decision-related evidence. In particular, we demonstrated that single-trial amplitudes of a late EEG component indexing decision evidence (Philiastides et al., 2006b; Philiastides & Sajda, 2006; Ratcliff et al., 2009b; Lou, Li, Philiastides, & Sajda, 2014) were amplified in the course of learning, such that these representations became more robust to noise (rather than a reduction in noise as such). In contrast, a temporally earlier component encoding sensory (stimulus) evidence - even in the absence of a face/car decision task (Philiastides et al., 2006b) - was not affected by training. These findings suggest that it is the strengthening of the connections between early sensory encoding and downstream decision-related processing that are driving perceptual learning in our task.

Crucially, we also showed that the onset of the late component (which on average coincides with the onset of decision evidence accumulation (Kelly & O’Connell, 2013; O’Connell et al., 2012; Philiastides, Heekeren, & Sajda, 2014) systematically moves earlier in time with training. This finding is particularly interesting since we have previously observed comparable temporal shifts in this component while manipulating task/stimulus difficulty (Philiastides et al., 2006b; Philiastides & Sajda, 2006; Ratcliff et al., 2009b). We view this as additional evidence that our learning effects on the late component lead to changes in perceptual sensitivity. More specifically, the earlier the onset time of the late component, the stronger the behavioural improvements. In other words, speed and accuracy improve with each day of learning, see Figure 2.2b-c, and the late component occurs earlier each day, see Figure 2.3d), consistent with a decrease in perceived task difficulty.

These temporal changes are also in line with a faster and more efficient accumulation of evidence as often predicted by sequential sampling models of decision making (Ratcliff & Smith, 2010; Ratcliff, Smith, & McKoon, 2015; Smith & Ratcliff, 2004) (for example, increases in the drift-rate and decrease in nondecision time variability). We return to this topic, and model these accumulation of evidence effects directly, in Chapter 5.

Research on perceptual learning has recently focused on the extent to which perceptual learning is due to improvements in sensory abilities that are (informationally and temporally) earlier than the decision process itself or due to improvements in post-sensory and decision-related processing. Consistent with the former account, several psychophysics studies have demonstrated that perceptual learning is often highly specific to the location and other properties of the stimuli (Ball & Sekuler, 1987; Crist et al., 1997; Fahle & Edelman, 1993; Fiorentini & Berardi, 1980; Karni & Sagi, 1991; Poggio et al., 1992; Sagi & Tanne, 1994) implying specificity to the trained retinal location (Fahle, 2004, 2005). Similarly human fMRI studies offered evidence of activity enhancements in retinotopic areas corresponding to the trained visual fields (Schwartz et al., 2002) and increased responses along the whole hierarchy of early visual areas that correlated with improvements in behavioural performance following training over the course of several weeks (Furmanski et al., 2004; Jehee, Ling, Swisher, van Bergen, & Tong, 2012). These results are further corroborated by EEG recordings in humans showing post-training improvements in early visually-evoked components over occipital electrode sites (Bao, Yang, Rios, He, & Engel, 2010; Pourtois, Rauss, Vuilleumier, & Schwartz, 2008; Censor, Bonneh, Arieli, & Sagi, 2009) and electrophysiological recordings in NHPs linking behavioural performance with improvements in perceptual sensitivity in primary sensory areas (Ghose et al., 2002; Yan et al., 2014).

In contrast, other psychophysical studies proposed that perceptual learning can also arise from changes in how sensory signals are read out or interpreted by decision-making mechanisms (Petrov, Doshier, & Lu, 2005; Doshier & Lu, 1999; Lu, Liu, & Doshier, 2010) rather than from changes in primary sensory areas as such. Neural evidence in support of this interpretation comes from electrophysiology studies on NHPs (Law & Gold, 2008b, 2009) demonstrating that perceptual learning on a motion discrimination task affects downstream decision accumulator areas, rather than regions encoding the sensory

evidence (that is, motion direction). Specifically, accumulator neurons improved responsiveness to the decision evidence in the course of learning (as reflected in steeper evidence accumulation slopes), with these improvements being proportional to the animals' performance on the task. Correspondingly, recent fMRI studies in humans started to explore the effect of learning on the activity and connectivity patterns of higher-level ventral temporal (Kuai et al., 2013; S. Li et al., 2009) and decision-related regions (Baldassarre et al., 2012; Chen et al., 2015; Kahnt et al., 2011; Lewis, Baldassarre, Committeri, Romani, & Corbetta, 2009).

These seemingly discrepant accounts of the temporal locus of perceptual learning may be reconciled by considering differences in the experimental demands of the task at hand. For example, a recent theoretical account proposed a unified two-stage model of perceptual learning (Shibata, Sagi, & Watanabe, 2014). According to this model, there are two distinct types of plasticity underlying perceptual learning: feature-based plasticity and task-based plasticity. On the one hand, feature based plasticity affects early sensory processing stages and occurs with mere exposure to stimuli, regardless of whether the stimuli are relevant to the task or not. Task-based plasticity, on the other hand, can be thought of as a higher-level processing stage arising from direct and active involvement in a behavioural task. In this formulation, the relative contribution of the two plasticity types to the overall enhancement in performance hinges largely on the training procedures, the stimuli and the intricacies of the task used in learning (W. Li, Piëch, & Gilbert, 2004).

More specifically, a distinction could be drawn between tasks that involve learning of relatively primitive stimulus features such as orientation, spatial frequency or contrast and those employing more complex stimuli such as objects and faces (T. Watanabe & Sasaki, 2015). Although learning of highly primitive features could occur locally at the level of early sensory processing, more complex stimuli (made up of a combination of primitive features) might require active involvement of downstream higher-level sensory or decision-related areas. In our design, for instance, complex object categories are used. In addition, phase discrimination, which is shown to involve processes beyond early visual cortex (Perna, Tosetti, Montanaro, & Morrone, 2008), is required to perform the task reliably. Our findings thus appear to rely heavily on the enhancement of the relevant stimulus representations during post-sensory, rather than early sensory processing.

In summary, our study provides insights into the neurobiology of perceptual learning and offers strong support to the notion that neuronal plasticity can occur at multiple time-scales and locations, depending on task demands and context. As such our findings can help revise existing theories of perceptual learning focusing only on early sensory processing and provide the foundation upon which future studies continue to interrogate the neural systems underlying perceptual decision making.

Chapter 3

A reinforcement learning account of the effect of learning on perceptual decision making

3.1 Reinforcement learning accounts of perceptual learning

As described in Chapter 1, RL (see Sutton & Barto, 1998, for general introduction) is a major subfield of machine learning that describes algorithms that adapt behaviours to maximize rewards and minimize punishments. RL has its origins in the study of classical conditioning and operant conditioning in behaviourist psychology and is still applied to the study of human and animal behaviour (see, for example, Wise & Rompre, 1989; Schultz, Dayan, & Montague, 1997; Schultz, 1998; O’Doherty, Dayan, Friston, Critchley, & Dolan, 2003; O’Doherty et al., 2004; Hyman, Malenka, & Nestler, 2006; Guitart-Masip et al., 2012; Huys, Guitart-Masip, Dolan, & Dayan, 2015; Huys, Maia, & Frank, 2016).

Law and Gold (2009) have applied RL to the study of perceptual learning. They argue that perceptual learning can be driven by a reinforcement signal that generates a selective readout of the sensory neurons that are most informative for any given perceptual decision. In Law and Gold (2008a), which we describe in Chapter 2 as the basis for the work we present there, Law and Gold (2008a) trained NHPS to decide the direction of motion

in a RDK and then respond with an eye movement saccade. In their study, Law and Gold (2008a) describe how the monkeys learn an association between the correct saccadic response and the motion of the RDK, even when the motion signal was relatively weak. As we describe in Chapter 2, the improved perceptual learning corresponded to changes in LIP neurons. The LIP area is known to involve decision making and reward processing. Law and Gold (2008a) propose that RL signals establish functional connections from sensory neurons in the MT area to a LIP decision neurons. These LIP neurons interpret the sensory information in order to determine the correct eye movement response. Law and Gold (2008a) further argue that during perceptual learning, RL signals continue to strengthen the connections between the most informative sensory neurons and the LIP area.

Law and Gold (2008a) model the above described process as follows. First, the MT area neurons are modelled as a set of $n = 7200$ neurons, each of which respond to a specific direction of motion uniformly distributed from 0° to 360° . The direction-of-motion signal of these n neurons at time t is represented by

$$x_1^t, x_2^t \dots x_i^t \dots x_n^t.$$

The LIP area accumulates these signals at time t into one cumulative variable v_t as follows:

$$v_t = \sum_{i=1}^n w_i^t x_i^t$$

where

$$x_i^t = \text{motion driven activity of neuron } i \text{ at trial } t$$

and

$$w_i^t = \text{weight given to neuron } i \text{ at trial } t.$$

If v_t is positive, this represents an average rightward motion. If v_t is negative, this represents an average leftward motion.

The weight vector

$$\mathbf{w}^t = w_1^t, w_2^t \dots w_i^t \dots w_n^t,$$

is initially, i.e. at time $t = 0$, random. The update of the weight vector from trial t to trial $t + 1$ is as follows:

$$\mathbf{w}^{t+1} = \mathbf{w}^t + \Delta \mathbf{w}^t.$$

In other words, we have

$$\begin{bmatrix} w_1^{t+1} \\ w_2^{t+1} \\ \vdots \\ w_n^{t+1} \end{bmatrix} = \begin{bmatrix} w_1^t \\ w_2^t \\ \vdots \\ w_n^t \end{bmatrix} + \begin{bmatrix} \Delta w_1^t \\ \Delta w_2^t \\ \vdots \\ \Delta w_n^t \end{bmatrix} \quad (3.1)$$

Here, $\Delta \mathbf{w}^t$ is based on difference between *expected* and *actual* reward. In particular,

$$\Delta \mathbf{w}^t = \alpha C_t \times (r_t - e_t) \mathbf{x}^t$$

where α is a learning rate, and C_t the directional response that is made at time t , i.e. $C_t \in \{-1, 1\}$, with $C_t = -1$ being a *left* response and $C_t = 1$ being a *right* response. Given that v_t is the motion-direction response at time t , with positive values of v_t indicating more rightwards movement and the negative values of v_t indicating more leftwards movement, then C_t is defined as follows:

$$C_t = \Phi(v_t > 0),$$

where $\Phi(\cdot)$ is a threshold function that takes on the value of 1 if its argument is positive and -1 if its argument is negative.

Crucially, on each trial, the monkey gets a reward r_t depending on whether its response is correct or not. If the response is correct, the reward is signified by $r_t = 1$ and if the response is incorrect, the reward is $r_t = 0$. Finally, the term e_t is the expectation of a reward at time t . This is given by

$$e_t = \frac{1}{1 + e^{-\beta |v_t|}}.$$

Note that as the absolute value of v_t increases e_t increases. Its rate of increase is controlled by the free parameter β . Therefore, e_t represents the confidence that a reward will be obtained and this rises when the cumulative signal v_t becomes increasingly positive or increasingly negative (signifying, respectively, stronger rightward and stronger leftward motion).

On each trial t , after each weight update, the weight vector is re-normalized by dividing each w_i^t by

$$w_i^t \leftarrow w_i^t / \sqrt{\frac{\sum_{i=1}^n (w_i^t)^2}{w_{\text{amp}}}}$$

where w_{amp} was a chosen constant. This guarantees that on all trials, the sum of the squared weight vector will always be equal to w_{amp} . The normalization occurs at each trial t and applies to the vector of $n = 7200$ weights at each trial t .

3.1.1 Extension of Law and Gold (2008a) by Kahnt et al. (2011)

Kahnt et al. (2011) have extended the work of Law and Gold (2008a) to make it applicable to work with humans and neuroimaging experiments. We illustrate this model in Figure 3.1. As with Law and Gold (2008a), Kahnt et al. (2011) represent the input stimulus at time t by x_t . In this case, however, x_t is a *scalar* rather than a *vector* quantity. It represents the orientation of the perceptual stimulus in a visual orientation discrimination task done by humans in an fMRI study. Analogously, to the v_t of Law and Gold (2008a), in Kahnt et al. (2011), the decision variable at time t

$$v_t = x_t \cdot w_t,$$

which is clearly a weighting of the input stimulus.

In Law and Gold (2008a), we saw that the behavioural response was a deterministic function of the decision variable v_t . In the case of Kahnt et al. (2011), the binary response is probabilistic. In particular, the probability of *right* is given by

$$p_t = \frac{1}{1 + e^{-\beta(v_t - c)}},$$

and so the probability of *left* is given by $1 - p_t$. The reward at time t , i.e., r_t takes on the value of 1 if the response is correct and zero if the response is incorrect. The *reward prediction error* at time t is

$$\delta_t = r_t - e_t,$$

where e_t is the expectation of a reward at time t and is given by

$$e_t = \frac{1}{1 + e^{-\beta|v_t - c|}}.$$

The weight variable is then updated as

$$w_{t+1} = w_t + \alpha \cdot \delta_t,$$

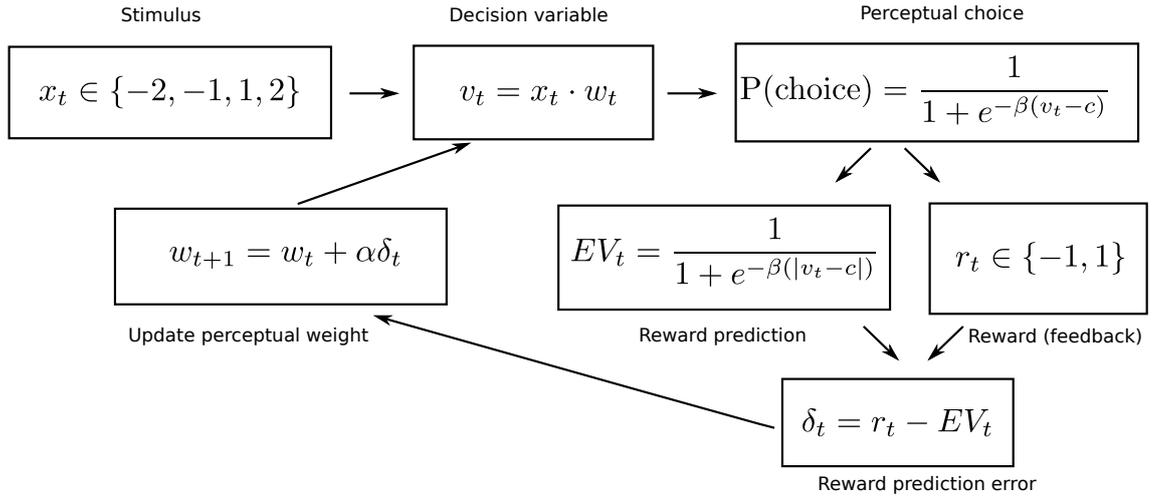


Figure 3.1: **Kahnt's reinforcement learning model of perceptual learning.** The subject's decision variable at time t , v_t , is a weighting of the sensory information, which in our case is $x_t \in \{-2, -1, 1, 2\}$, weighted by w_t . The DV_t is then converted into probability of a response. The higher this probability, the higher the probability that subject makes a "face" response. The lower this probability, the higher the probability that subject makes a "car" response. The model updates the weight on each iteration based on whether the subject is correct or not in their response.

where α is the learning rate. Note that in this model, there are 4 free parameters: the learning rate α , the sensory gain parameter β , the sensory offset parameter c , and the initial value of the weight, i.e. w_0 .

3.2 Modelling face versus car perceptual learning

Following Kahnt et al. (2011), we use a RL model to model how feedback leads to increased accuracy of responses over the trials in the learning task described in Chapter 2.

Assuming that the stimulus presented on trial t is denoted by x_t , then following the Kahnt et al. (2011) model, the probability of "face" response is

$$p_t = P(\text{response} = \text{face} | \text{trial} = t) = \frac{1}{1 + e^{-\beta \cdot (v_t - c)}}, \quad \text{where } v_t = x_t \cdot w_t.$$

In our case, the x_t stimulus variable codes high coherence faces, low coherence faces, low coherence cars, and high coherence cars as the values 2, 1, -1, 2, respectively¹. Starting

¹In coding our stimuli as such, we are making simplifying assumptions such as that the perceived difference between the low and high coherence faces is the same as the perceived difference between the low

with a constant value of the weight variable trial $t = 0$, our RL model updates the value of $w_1, w_2, w_3 \dots w_T$ exactly as in Kahnt et al. (2011):

$$w_{t+1} = w_t + \alpha \cdot \delta_t$$

where

$$\delta_t = r_t - e_t, \quad e_t = \frac{1}{1 + e^{-\beta \cdot |v_t - c|}}.$$

Here, r_t signifies the feedback that they receive on trial t , i.e. $r_t = 0$ if they are correct and $r_t = 1$ if they are incorrect. As with Kahnt et al. (2011), we have free parameters, α , β , c , w_0 .

Model performance was quantified in terms of the discrepancy between the behavioural accuracy of the subject and the model's accuracy, i.e.,

$$\Omega \triangleq |\text{Subject accuracy}_{\text{low}} - \text{Model accuracy}_{\text{low}}| + |\text{Subject accuracy}_{\text{high}} - \text{Model accuracy}_{\text{high}}|.$$

Here, the accuracies are calculated separately for the low coherence trials, i.e. where $x_t \in \{-1, 1\}$, and the high coherence trials, where $x_t \in \{-2, 2\}$, with these accuracies for the model calculated by

$$\text{Model accuracy}_{\text{low}} = \sum_{\{t: x_t=1\}} p_t + \sum_{\{t: x_t=-1\}} q_t,$$

$$\text{Model accuracy}_{\text{high}} = \sum_{\{t: x_t=2\}} p_t + \sum_{\{t: x_t=-2\}} q_t,$$

with p_t being the model's probability of a "face" response, defined above, and $q_t = 1 - p_t$ being the model's probability of a "car" response.

As in Kahnt et al. (2011), the variables α , β , c and w_0 were free parameters that had to be optimized for each model. This was accomplished by finding the values that minimized Ω by an exhaustive search over a grid of value ranges for each variables.

In particular, for each subject on each of the three days of learning, we searched α over a grid of $N = 30$ values from 0.001 to 2.00; we searched β over a grid of N values from 0.001 to 5.00; we searched c over a grid of N values from -2 to 2.00; we searched and high coherence cars, and also that all the face and car stimuli can be represented on a one-dimensional continuum. Although we acknowledge that these are simplifying assumptions, and as such, we must be mindful of these when drawing any substantive conclusions from the model, nonetheless in any computational modelling, simplifying assumptions are inevitable.

w_0 over a grid of N values from 0.001 to 5.00. We chose these ranges, and the number of values of these ranges, after extensive experimentation with different possibilities.

3.3 Results of modelling face versus car perceptual learning

The average (with standard error in brackets) of free parameters over subjects are as follows, where we also present the average error Ω per day:

day	α	β	c	w_0	Ω
1	0.34 (0.10)	3.93 (0.27)	-0.03 (0.02)	0.86 (0.07)	0.02 (<0.01)
2	0.62 (0.16)	3.58 (0.42)	-0.14 (0.12)	0.60 (0.09)	0.03 (0.01)
3	0.63 (0.13)	3.84 (0.35)	0.03 (0.04)	1.04 (0.32)	0.03 (0.01)

We found a highly significant correlation between the behavioural accuracy of each subject on each day of training and the average accuracy of the model ($r = 0.95, p < 0.01$), see Figure 3.2a. Also, in Figure 3.2, we show the relationship, for one representative subject over the three days of learning, between the RL model’s decision variable (Figure 3.2b) and the time course of the early temporal component’s y variable (Figure 3.2c) and the late temporal component’s y variable (Figure 3.2d). From these plots, we note that there is an apparently stronger resemblance between the RL model’s decision variable and the late y variable than with the early y variable.

3.3.1 Analysis of the y variable

The primary focus of our analysis was to determine if there was a relationship between the decision variable in the RL model and the y decision variable that was obtained from the single trial EEG analysis. In Figure 3.3, we provide a plot of the relationship between the RL model’s decision variable and the early and late component values for each subject.

In order to assess whether there was a closer correspondence between the RL model’s decision variable and the early or late y variable, we perform a model comparison. The following multilevel linear regression of how the v_t varies as a function of y_t^{early} was used:

$$v_t = \alpha_{h_t} + \beta_{h_t} y_t^{\text{early}} + \epsilon_t$$

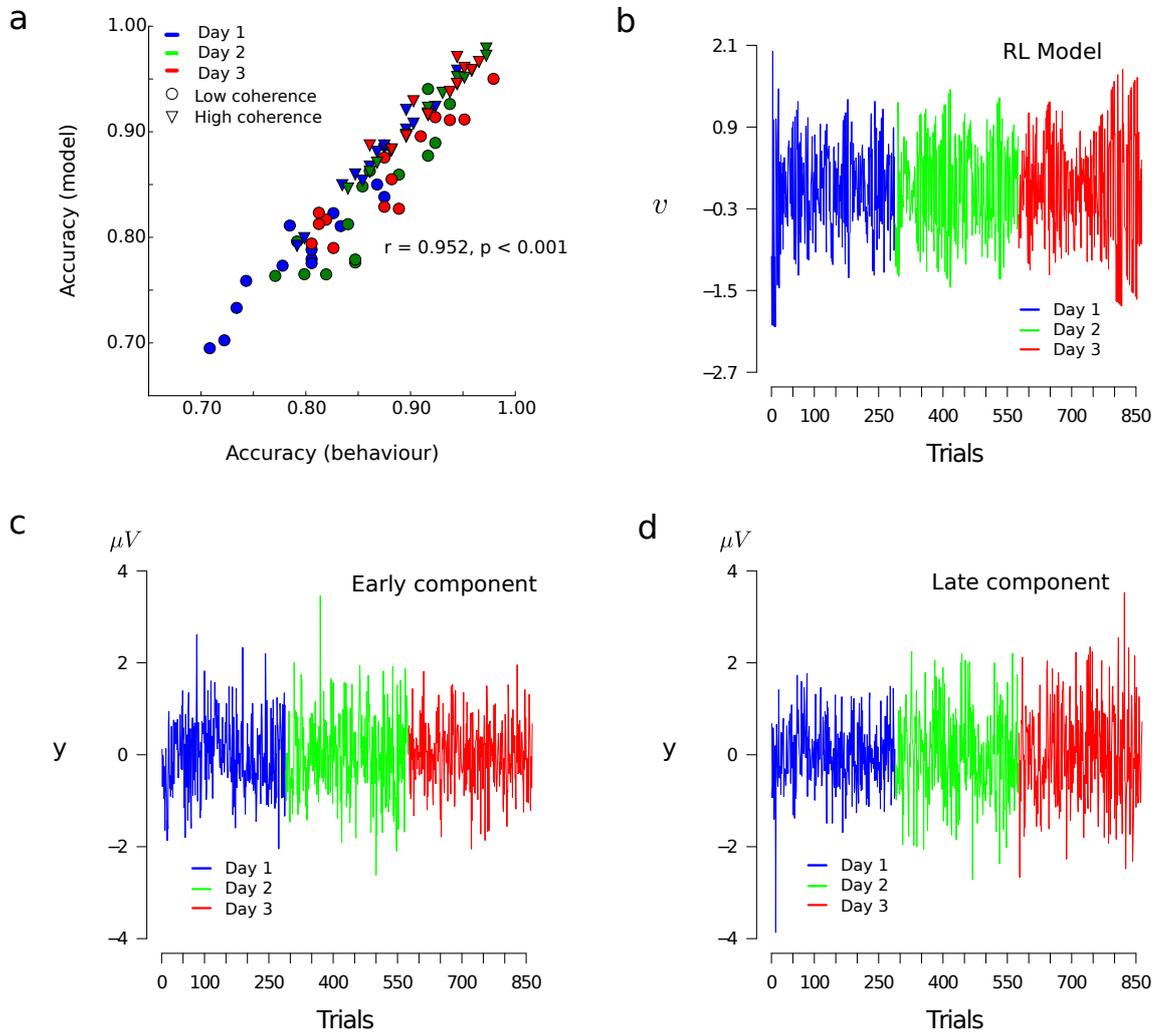


Figure 3.2: **Reinforcement learning model for perceptual choices.** (a) Scatter plot showing the correlation between the performance of all subjects and models, over the three training days and the two levels of stimulus phase coherence (low and high). (b) An example of trial-wise values of the decision variable from the RL model over the three days (from blue, green to red) for one subject (288 trials per day). (c) Individual y variable from the early temporal component for a representative subject over the course of the three training days (from blue, green to red), with 288 trials per day. (d) Individual y variable from the late temporal component for a representative subject over the course of the three training days (from blue, green to red), 288 trials per day.

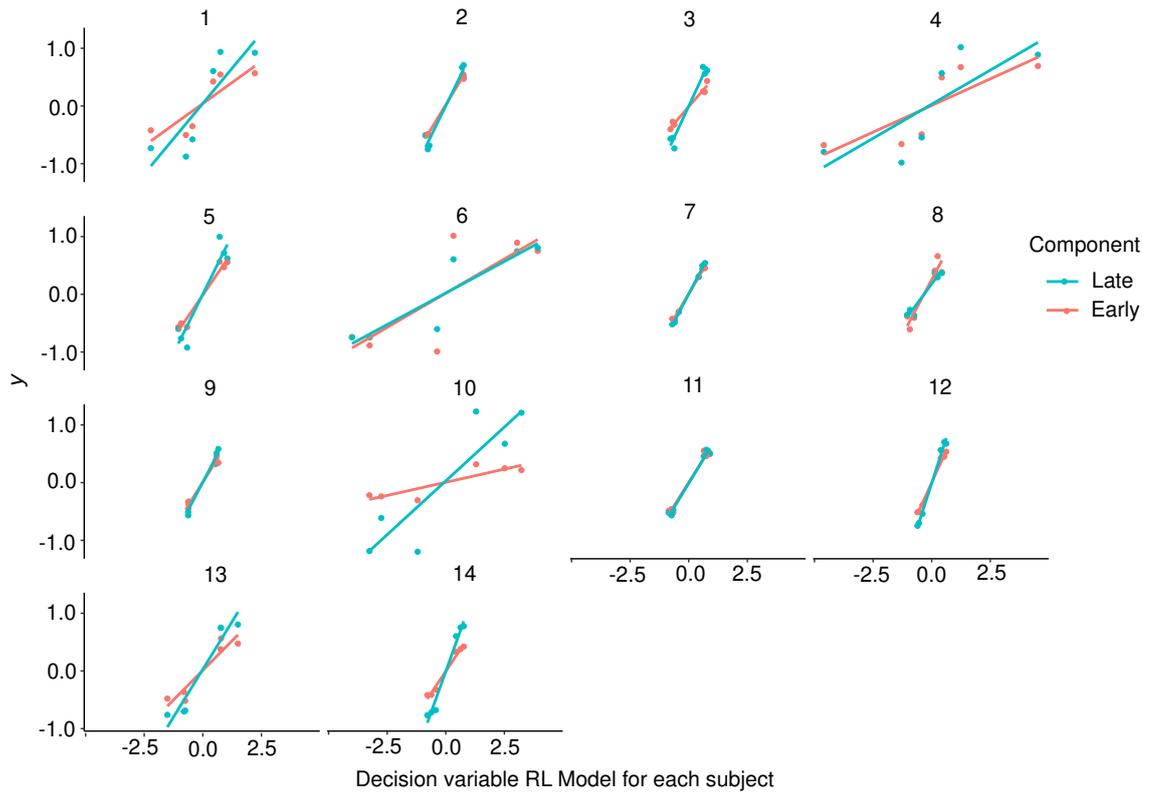


Figure 3.3: **Reinforcement learning model for perceptual choices for each subject.** Correlation between the RL model's decision variable with the y variable for the early temporal component and for the late temporal component, for all 14 subjects, averaged over all trail.

where

$$\alpha_j \sim N(a, \tau_a^2), \quad \beta_j \sim N(b, \tau_b^2)$$

are random effects on the intercept and slope that account for inter-subject variability in the model. A similar multilevel linear regression of how the decision variable varies as a function of y_t^{late} was then estimated:

$$v_t = \alpha_{h_t} + \beta_{h_t} y_t^{\text{late}} + \varepsilon_t$$

where the random effects are defined as above. Fitting these two models using R (lme4 package), we obtain

- BIC for early model: 51652.01.
- BIC for late model: 51525.85.

As before, we calculate the log Bayes factor as

$$-(\text{BIC}_{\text{late}} - \text{BIC}_{\text{early}})/2 = 63.07$$

and so $2 \times \log(\text{Bayes Factor}) = 126.15$. As noted on in Section 2.4.4 (page 47), using the widely used rules presented in Kass and Raftery (1995), if $2 \times \log(\text{Bayes Factor}) > 10$, this is *very strong evidence* in favour of the late component being a better predictor than the early component of the RL model’s decision variable.

3.4 Extensions

The model and analysis we have presented here have been extended by F. Queirazza, M. Philiastides and Diaz, J., and the results were presented in Diaz et al. (2017). In what follows, we describe this model following the description in Diaz et al. (2017)².

In the new model, as was the case above, perceptual decisions were driven by a decision variable denoting the subject’s hidden representations of the association between sensory evidence and stimulus category. The strength of the decision variable was modulated by dynamic updates of category-specific perceptual weights based on feedback information, thereby accounting for potential differences in learning trajectories between the stimulus categories. Compared with previous work just described that used a single stimulus-invariant perceptual weight (Law & Gold, 2009; Kahnt et al., 2011), the introduction of category-specific perceptual weights was designed to capture subject-wise choice biases, in that subjects might have a choice bias towards cars or faces and likewise might display an increasing ability to recognize cars or faces throughout the task.

In addition, the perceptual weights comprised signal and noise weights. While the former were designated to enhance stimulus representations in the course of learning, the latter accounted for the interference exerted by the antagonistic stimulus against the acquisition of the correct sensory associations. Thus, in this model, perceptual learning is expected to occur through gradually increasing signal weights as well as gradually decreasing noise weights. Compared with previous reinforcement-learning-like perceptual models (Law & Gold, 2009; Kahnt et al., 2011), this better captures instances whereby

²The purpose of presenting these additional analyses is because they based on some different modelling details, and as such they may provide some valuable new perspectives on the analyses just described.

improved task performance depends both on greater ability to recognize a given stimulus and on greater ability to rule out the antagonistic stimulus. In other words, on a face trial, subjects might correctly choose face partly because they are able to identify face-like features and partly because they are able to recognize that there are no car-like features.

As perceptual learning progresses, the estimates of signal and noise weights grow apart. As a result, the readout of sensory evidence is increasingly enhanced, reflecting the improving ability to discriminate between perceptual stimuli in the course of training.

This model was fitted to individual participant data and found a highly significant correspondence between the model's accuracy predictions and actual behaviour ($r=0.88$, $p < 0.01$; see Figure 3.4 a). Consistent with an enhanced readout of sensory evidence, a subject-wise gradual build-up in the trial-by-trial estimates of the decision variables was observed (see Figure 3.4 b). Correspondingly, a gradual increase in the model's signal weights were observed (see Figure 3.4 c). Between-day comparisons (1 versus 2, and 2 versus 3) of subject-wise mean decision variables (paired t-test: 1 vs 2 $t(13)=6.77$, $p < 0.01$; 2 vs 3 $t(13)=-2.36$, $p=0.02$; see Figure 3.4 d, e), and aggregate perceptual weights (signal weights: paired t-test: 1 vs 2 $t(13)=-6.74$, $p < 0.01$; 2 vs 3 $t(13)=-2.36$, $p=0.02$; noise weights: paired t-test: 1 vs 2 $t(13)=6.74$, $p < 0.01$; 2 vs 3 $t(13)=2.35$, $p=0.02$) revealed a significant effect of learning as observed in behaviour.

To offer neurobiological validity to the model, the single-trial decision variables estimated by the model were compared with the EEG component amplitudes that were obtained in the accompanying EEG experiments, described in Chapter 2. It was predicted that if the brain computes a version of the model-based decision variables to drive choices then one should observe a systematic amplification of the decision variable with training and a significant correlation with our late EEG component shown to index decision evidence. To this end, another regression analysis was run whereby the single-trial amplitudes of our early and late components were used to predict the model's decision variables. It was found that our late component was both a reliable predictor of the model's decision variables (Figure 3.4 f; t-test, $t(13)=21.81$, $p < 0.01$) and a significantly better predictor than the early component (Figure 3.4 f; paired t-test, $t(13)=3.06$, $p=0.01$).

Second, trials were separated into four bins (quartiles) based on the model-predicted magnitudes of the prediction error (PE) signal, which is thought to guide learning. A

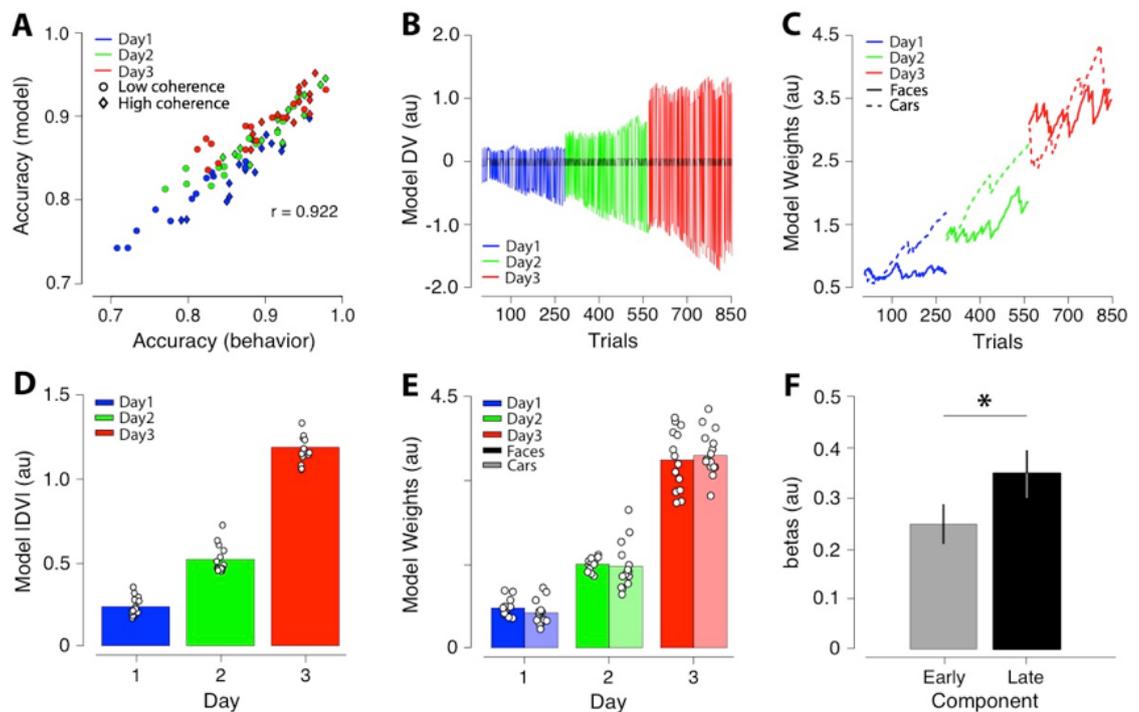


Figure 3.4: **Extended reinforcement learning model for perceptual choices.** (a) Scatter plot showing the correlation between the performance of individual subjects and models, over the three training days and the two levels of stimulus phase coherence (using the winning model). (b) Individual trial estimates of the model’s decision variable for a representative subject over the course of the three training days, superimposed on the amount of stimulus-defined sensory evidence (black trace). (c) Signal (positive) and noise (negative) perceptual weights for faces (solid lines) and cars (dashed lines) over the three training days for the same subject shown in b. (d) Average magnitude of the model’s decision variables across subjects over the course of the three training days. Individual subject data are also shown as point estimates. (e) Average signal (positive) and noise (negative) perceptual weights for faces (brightly coloured bars, left) and cars (faintly coloured bars, right) over the three training days. Individual subject data are also shown as point estimates. (f) Average regression coefficients (betas) reflecting the trial-by-trial association between the model’s decision variables and the amplitudes of the early and late EEG components estimated over all training days. Individual subject data are also shown as point estimates.

single-trial discriminant analysis was then run on feedback-locked EEG data between the very low and very high PE trial groups (that is, the middle two quartiles as were kept as 'test' data). This analysis revealed a centroparietal EEG component peaking on average at 354 ms post-feedback (Figure 3.5a). The timing and topography of this component are consistent with previous work on feedback-related processing in the human brain using a probabilistic reversal learning task (Fouragnan, Retzler, Mullinger, & Philiastides, 2015; Philiastides, Biele, Vavatzanidis, Kazzer, & Heekeren, 2010). To test whether this EEG component was parametrically modulated by the magnitude of the PE signal, discriminator amplitudes (y) were computed for trials with intermediate magnitude levels (those left out from the original discrimination analysis). Specifically, the spatial filter of the window that resulted in the highest discrimination performance for the extreme PE magnitude levels was applied to the EEG data with intermediate values. It was expected that these 'unseen' trials to show a parametric response profile such that the resulting mean component amplitude at the time of peak discrimination would proceed from very low < low < high < very high PE magnitude. Using this approach, it was demonstrated that the mean discriminator output for each quartile increased as a function of the model's PE magnitude (all pair-wise t-test comparisons across adjacent trial groups: p values <0.01; Figure 3.5b), thereby establishing a concrete link between the model's PE estimates and the feedback-related EEG component.

Taken together, these findings provide further evidence that perceptual learning enhances decision-related evidence, probably through a reinforcement-learning-like mechanism.

3.5 Discussion

The RL model of the behavioural data has implications for our understanding of the relationship between general reward based learning and decision making mechanisms in perceptual learning. Law and Gold (2009) show that the results reported in Law and Gold (2008a) can be explained in terms of RL. They modelled the connections between the sensory neurons in the MT area and the decision making centre of LIP, and used a RL signal to modify the connections between these two areas. This explanation was corroborated

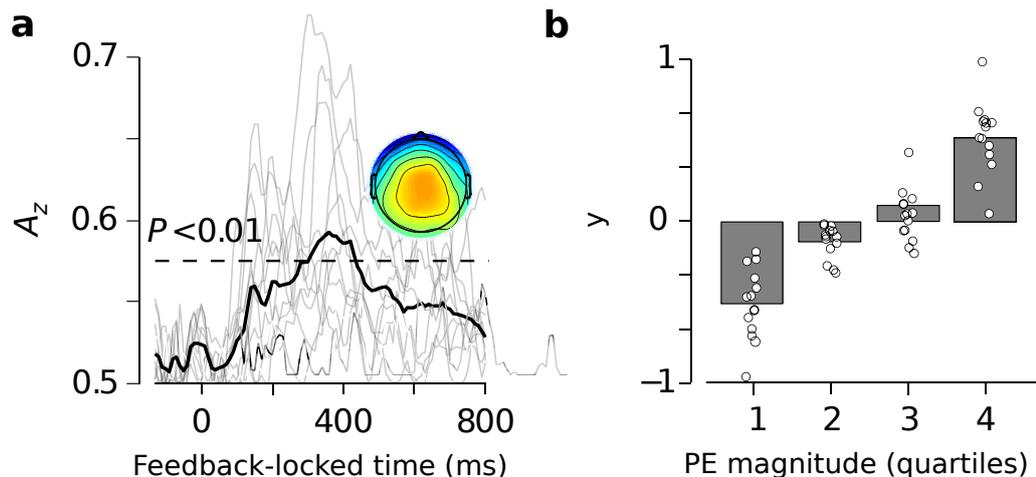


Figure 3.5: **Electrophysiological correlates of prediction error (PE).** (a) Multivariate discriminator performance (A_z) during very low versus very high PE magnitude trials on feedback-locked EEG responses averaged across subjects and days revealing a late PE component. Discriminator performance and component peak times were comparable across the three days. The dotted line represents the average A_z value leading to a significance level of $p = 0.01$, estimated using a bootstrap test. Faint lines represent individual subject data. Inset: average scalp topography associated with the PE component, estimated at the time of subject-specific maximum discrimination. (b) Mean discriminator amplitude (y) for the PE component, binned in four quartiles based on model-based estimates of the magnitude of the PE, showing a clear parametric response along the four trial groups. Quartiles 1 and 4 were used to train the classifier, while quartiles 2 and 3 contain ‘unseen’ data with intermediate PE magnitude levels. Individual subject data are also shown as point estimates.

by the work of Kahnt et al. (2011) who trained subjects on in a visual discrimination task where feedback was given. Their behavioural results are well modelled by a RL model that changes how the sensory information is interpreted.

Consistent with previous accounts (Law & Gold, 2009; Kahnt et al., 2011) we also showed that the learning-induced behavioural improvements in the visual discrimination task that we presented in in Chapter 2 and Diaz et al. (2017) could be reliably explained in terms of a RL mechanism. More specifically, we showed that a model that uses a prediction error signal (Rushworth, Mars, & Summerfield, 2009; Schultz et al., 1997; O’Doherty et al., 2003; Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006) to continuously adjust the stimulus specific perceptual weights on the sensory evidence (Guggenmos, Wilbertz,

Hebart, & Sterzer, 2016) led to amplification of the relevant stimulus representations in the course of training (making them more robust to noise). We further demonstrated that trial-by-trial changes in our late EEG component shown to index decision evidence reliably tracked the amplification of sensory information predicted by the model. These results imply that perceptual learning involves an enhanced readout of sensory information during decision making likely via a RL-like process, endorsing the view of a domain-general learning mechanism (Rushworth et al., 2009). Although it is true that our task did not involve any explicit reward as a reinforcer, we view the implicit rewarding nature of correct response as a 'teaching signal' for strengthening the neural representation of sensory contingencies (Guggenmos et al., 2016).

Chapter 4

The neural signatures of the role of prior expectation in perceptual decision making

4.1 Introduction

In natural environments, perceptual stimuli do not occur in isolation or independently of their temporal and spatial context. As such, prior expectation — or the expectation of the properties or identity of a stimulus based on its temporal or spatial context — will affect the speed and accuracy of perceptual decision making. This has been repeatedly demonstrated behaviourally (e.g., Bar, 2004; Oliva & Torralba, 2007; Fiser & Aslin, 2002). However, as we have described in Chapter 1, the neurobiological mechanisms underlying how expectations affect perceptual decision making are not well understood. In particular, there is an open question concerning whether prior expectation affects the baseline activity or the evidence accumulation in an accumulation-to-bound decision making system, i.e. one that could be modelled by a sequential sampling model. As we have described in more detail in Chapter 1, studies with human and NHPs done by Basso and Wurtz (1998); de Lange et al. (2013); Albright (2012); Puri et al. (2009) point to the locus of prior expectations being related to changes in baseline activity of the decision making systems. However, in a study using NHPs by Hanks et al. (2011), it was shown that the rate of firing of LIP neurons — known to integrate sensory signals towards a decision threshold

— increased with increasing match between the prior expectation and the actual observed stimuli. EEG results with human subjects by Cravo et al. (2013) also point to the role of prior expectation in affecting the efficiency of information accumulation, specifically by modulating the signal-to-noise gain of visual processing. Behaviourally consistent results with human subjects have been reported by Dunovan et al. (2014), who demonstrated that the primary role of prior expectation is to bias the rate of evidence accumulation favouring the more probable stimulus category.

On the basis of the work described in for example Piliastides and Sajda (2006); Piliastides et al. (2006a); Ratcliff et al. (2009b); Piliastides et al. (2006a), and from our work in Diaz et al. (2017) and in Chapter 2 of this thesis, we argue that there are two main temporal EEG components in perceptual decision making. The early component is involved in sensory encoding, while the late component is the neural signature of the post-sensory, decision-relevant evidence. As such, and if prior expectation affects information accumulation for decision making, in a perceptual decision making task, we should observe the late temporal component, but not the early one, being affected by the prior probability of the stimulus as indicated by a pre-stimulus cue. Importantly, we also expect an interaction between prior probability of the stimulus and learning. In particular, we predict that as the role of cues become more fully learned, the late temporal component, but not the early temporal component, will vary with the increasing match between the prior expectation and the actual observed stimulus. In other words, on those trials where the pre-stimulus cue predicts the stimulus that does in fact occur (e.g., when the cue predicts a face and a face image occurs, or when the cue predicts a car and a car image occurs), we should see stronger values in the late temporal component compared to those trials where the pre-stimulus cue predicts a stimulus that does not then occur (e.g. when the cue predicts a face but a car image occurs, or when the cue predicts a car but a face image occurs). By contrast, we do not predict any such effect on the early component.

Therefore, the aim of the present research is to investigate the neural signatures of the effect of prior expectation in perceptual decision making. This will be accomplished using a combination of single-trial EEG analysis, regression modelling of speed and accuracy response data, and also diffusion modelling of these behavioural results (following Dunovan et al., 2014). In this chapter, we will present the results of the EEG analysis and

speed and accuracy regression modelling. In Chapter 5, on the other hand, we provide the full details of the diffusion modelling.

In the behavioural task, we will follow an identical paradigm used in our previous studies: Subjects will perform a perceptual discrimination task, i.e., face/car discrimination, over the course of three days; two different levels of stimulus noise will be used. However, in addition, at the start of each trial one of three different cues (see below for details) will be presented. Each of these cues has an associated probability of the upcoming stimulus being a face or a car.

In terms of the behavioural responses, we predict a decrease in reaction time and an increase in accuracy with the increasing match between prior expectation and the observed stimulus. We also predict that the effects just mentioned will emerge gradually over the course of the three days. In other words, they will be relatively weak on the first day, stronger by the second, and strongest by the third. Finally, we predict that there will be an interaction between the effects of cue probability and training day. In particular, the effect of the prior expectation will be *decrease* with increased training. This particular prediction derives from Bayesian principle that prior information should influence decision making more when most when sensory evidence is weaker (e.g. Bogacz et al., 2006). More importantly, however, on the basis of the EEG single-trial analysis, we predict that when these cues are fully learned, consistent with Hanks et al. (2011), the strength and variability of the late temporal component, but not the early temporal component, will vary with the increasing match between the prior expectation and the actual observed stimulus.

4.2 Materials and Methods

4.2.1 Subjects

16 Subjects (5 men and 11 women, age range 21-35 years) participated in this study. Each subject performed the experiment on three consecutive days. This entailed that there were a total of 48 EEG testing sessions. All participants were right handed, reported normal vision and no history of neurological problems. Informed consent was obtained

from all participants in accordance with the guidelines of the Institute of Neuroscience and Psychology from Glasgow University.

4.2.2 Stimuli

The stimuli were selected from the same stimulus set described in e.g., Piliastides and Sajda (2006, 2007) and that were used in the experiments described in Chapter 2 and in the pilot experiments of this chapter. As such, we will not repeat their details here, but see Section 2.2.2 on Page 25.

4.2.3 Behavioural Task

The behavioural task is shown in Figure 4.1a. Subjects were presented with the stimuli just described and performed a two alternative forced choice classification task whereby they classified each image as either a face or a car. Subjects sat a distance of 75 cm from the computer monitor. At the start of each trial, a text-based cue was displayed for a duration of 750 ms. There were three different cues: 30C/70F, 50C/50F, 70C/30F. Each one stated the probability that the upcoming stimulus would be either a Face or a Car. For example, 30C/70F indicates that the probability of Car is 30% and the probability of a Face is 70%. After the cue, a blank screen was displayed for a random duration that ranged uniformly between 1.0 to 1.5 seconds. The stimulus image was then presented for 50 milliseconds and subjects were given up to 1250 milliseconds to make their classification response, which was done using a USB button box using their right hand's index (for face response) and middle (for car response) fingers. No feedback was given about whether the response was correct or incorrect. The trials were presented in 5 blocks of 72 trials, with a 60 second rest period between each block. The entire experiment lasted approximately 25 minutes. Each subject performed this task on three consecutive days. With the experiment taking place at the same time on each day so that there was 24 hours between each session for all subjects. On the first day, subjects performed a practice session of the face/car classification task but with a different set of face and car images.

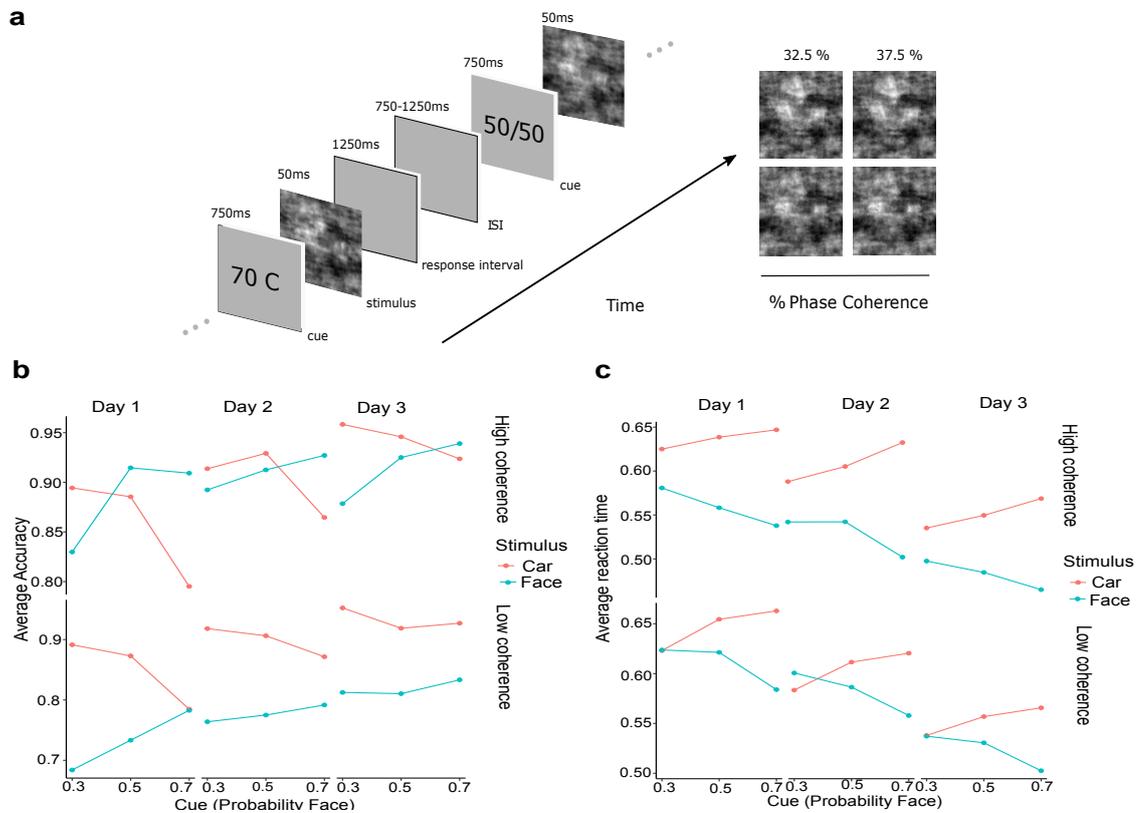


Figure 4.1: Experimental design and behavioural analysis. (a) Schematic representation of the experimental paradigm. The behavioural task in the EEG experiment involved verifying if noisy stimuli were either faces or cars. Subjects had to categorize a noisy image presented for 50 ms as a face or a car and indicate their choice with a button press within 1,250 ms following the stimulus presentation. Prior to stimulus onset, a cue was shown that had a fixed probability of being followed by a face or car stimulus, followed by an inter-stimulus interval that varied randomly between 1 and 1.5s. Subjects performed this task on three consecutive training days. A sample face image (upper row) and car image (lower row) at the two levels of phase coherence used in the task (32.5% and 37.5%) are shown on the right. (b) Average accuracy as a function of pre-stimulus cue for the two stimulus types (face: blue; car: red), two coherence levels (high & low) over the three days of learning across all 16 subjects. (c) Average reaction time for accurate responses as a function of pre-stimulus cue for the two stimulus types (face: blue; car: red), coherence level (high & low), and three training days across all 16 subjects.

4.2.4 EEG Data Acquisition

Subjects performed the task on three consecutive days, in a dark and soundproof room. As they performed the task, their EEG was recorded with a 64 channel Ag/AgCl scalp electrode actiCAP EEG system (Brain Products GmnH, Gilching, Germany). The active ground electrode was placed just below the Pz electrode of the International 10-20 system method. The active reference electrode was placed on the left mastoid. The impedance was always below 5kOhm for each subject on each day. The EEG signal was acquired in an identical manner to that described in Section 2.2.4 on Page 27.

4.2.5 Single Trial EEG Data Analysis

We performed single trial discriminant analysis in order to describe how temporal dynamics of decision making changes as a consequence of perceptual learning. Here, we followed the identical procedure described in the experiments in Chapter 2, which was based closely on the paradigm that has been established in previous studies by Philiastides et al. to identify the neural signatures of perceptual decisions. Given that we have described this procedure in depth above, we will not repeat its details here, but see Section 2.2.5 on Page 27.

4.3 Results

4.3.1 Behavioural analysis

In the analysis of the behavioural data, see Figure 4.1b-c (page 72), the model fit measure that we will use here is the Akaike Information Criterion (AIC). AIC is defined as follows:

$$-2\log L + 2Q,$$

where L is the likelihood of the model at its maximum likelihood estimator and Q is the number of parameters in the model. AIC can be seen as an approximation to cross-validation (see, e.g., Fang, 2011). Cross-validation is a measure of out-of-sample model generalization that *holds-out* a sub-sample of the data and assesses how well a fitted model can predict the held-out data. Usually, multiple alternative held-out data-sets are used,

with each one being predicted by the model fitted to the not held-out data. The overall predictive accuracy measure is then based on the average over the predicted accuracy calculated on each alternative held-out data-set. Being an approximation to cross-validation, AIC can be used as a measure of how well any given model will generalize to new data that is assumed to have been generated by the same probabilistic generative process that generated the original data set . Like other so-called *information criteria* such as BIC or Deviance Information Criterion (DIC), the higher the value of AIC, the *poorer* the model fit measure. In other words, lower values of AIC indicates better model fit.

In all analyses, whether of the accuracy or reaction time, there are 4 main predictor variables: stimulus type, cue probability, coherence level, and day of learning. There are potentially many interaction between these predictors: one four-way interaction, 4 three-way interactions, 6 two-way interactions, and 4 main effects. Some or many of these interactions may be redundant. In order to identify a minimal model without redundant predictors, we use a backward elimination process. Starting with the full model, i.e. the model with all possible interactions, we calculate its AIC value, and then drop higher-order interaction terms if doing so does not lead to any substantial increase in the AIC value. We then proceed to drop the next highest interaction terms and so on until no change in AIC occurs when any additional term is dropped. Note that an increase in AIC of 4 or more is required for the dropped term to be seen as non-redundant. In this, we are following the general advice of Burnham and Anderson (2003).

Accuracy analysis

In Figure 4.1b, we plot the average accuracy per cue for high and low coherence trials and for faces and cars.

In the following table, we show the AIC of the full logistic regression model for accuracy, and the same model with the four-way interaction term dropped.

Full model	Drop 4-way interaction
10984	10983

Clearly, dropping the 4-way interaction does not lead to an increase in the AIC value and so the model without the 4-way interaction is preferred. We now drop the three-way terms individually and calculate the AIC values of the resulting models:

Full three-way model	Drop stim x prob x coh	Drop stim x prob x day
10983	10981	10983

Drop stim x coh x day	Drop prob x coh x day
10982	10981

From this, we see that dropping all three-way interaction terms also does not lead to any increase in AIC, and so all of these terms can be treated as redundant and dropped with loss of model effectiveness.

Full two-way model	Drop stim x prob	Drop stim x coh	Drop stim x day
10980	11053	11048	10994

Drop prob x coh	Drop prob x day	Drop coh x day
10979	10978	10978

Here, we see that the pairwise interaction between stimulus and probability, stimulus and coherence, stimulus and day all result in a notable increase in AIC and as such need to be preserved in the model. On the other hand, we may drop, without loss, the pairwise interaction between probability and coherence, probability and day of learning, and coherence and day of learning.

Proceeding with this reduced model, the coefficients for the fixed effect terms are as follows:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	2.43	0.26	9.39	< 0.01
stimulusface	-0.99	0.21	-4.64	< 0.01

probability	-1.59	0.25	-6.41	< 0.01
coherencelow	-0.14	0.08	-1.80	0.07
day	0.50	0.05	10.04	< 0.01
stimulusface:probability	2.81	0.32	8.70	< 0.01
stimulusface:coherencelow	-0.91	0.10	-8.65	< 0.01
stimulusface:day	-0.26	0.06	-4.13	< 0.01

This analysis reveals the effect of cue probability on accuracy. There is a significant increase in the probability of an accurate response for the car stimuli as the probability of car stimuli on the upcoming trials, as indicated by the cue, increases, i.e. coefficient for the “probability” coefficient (with car being the base category for the stimulus variable) is -1.59, $p < 0.01$. Note that the negative slope shows that as the probability of the upcoming stimulus being a face increases, which means the probability of the upcoming stimulus being a car decreases, the accuracy of recognizing a car decreases. Likewise, there is a significant change in this effect of the cue probability for the face stimuli. In particular, the slope changes from -1.59 to $(-1.59 + 2.81) = 1.22$, and this change is significant, $p < 0.01$. Also of note is that there is an increase in accuracy with each day of learning. This occurs for car stimuli, with the coefficient being 0.50, $p < 0.01$. For face stimuli, the coefficient changes to $(0.50 + -0.26) = 0.23$, and the 95% confidence interval for this new slope is from 0.11 to 0.36. There is also a general decrease in accuracy for face stimuli, with the coefficient being -0.99, $p < 0.01$.

Reaction time analysis

In Figure 4.1c, we show the average reaction time as a function of our three predictor variables, using only the data from accurate trials.

Given that 88% of trials were answered accurately, we will focus our analysis on only the accurate trials. As for the accuracy analysis, we begin our analysis by dropping redundant terms. For this, we will take advantage of the `step` command in the `lmerTest` R package that allows for automatic backward elimination of terms in linear random effects models (Note that this was not an option for use with the random effects logistic regression for the accuracy analysis above).

Below, we show the coefficients of the optimal model that results from this elimination procedure:

	Estimate	Std. Error	df	t value	Pr(> t)
(Intercept)	-0.47	0.03	18.81	-13.63	< 0.01
stimulusface	0.05	0.01	15114	4.57	< 0.01
probability	0.16	0.01	15113	10.56	<0.01
coherencelow	0.01	<0.01	15124	1.51	0.13
day	-0.08	<0.01	15108	-37.60	<0.01
stimulusface:probability	-0.36	0.02	15110	-17.01	<0.01
stimulusface:coherencelow	0.08	0.01	15120	11.41	<0.01

These results closely parallel those of the accuracy analysis. This analysis clearly shows the effect of the cue probability on reaction time. There is a significant decrease in reaction time for the car stimuli as the probability of car stimuli on the upcoming trials, as indicated by the cue, increases, i.e. coefficient for the “probability” coefficient (again “car” is the base category for the stimulus variable) is 0.16, $p < 0.01$. Likewise, there is a significant change in this effect of the cue probability for the face stimuli. In particular, the slope changes from 0.16 to $(0.16 + -0.36) = -0.21$, and this change is significant, $p < 0.01$. As before, there is an average decrease in reaction with each day of learning. This occurs for both face and car stimuli, with the coefficient for car being -0.07, $p < 0.01$, and there is no change to this effect for face stimuli. We note also that there is a significant increase reaction time with low coherence images for face stimuli, with the interaction effect being 0.07, $p < 0.01$.

4.3.2 EEG Analysis

We carried out a face versus car, stimulus onset locked, discrimination analysis. This is the identical kind of discrimination analysis that was done for the first project, i.e. as described in Chapter 2. The A_z value that shows how well the EEG signal discriminates between face and car stimuli at each millisecond after the onset of the stimulus is show in Figure 4.2a. Note, the darker lines show the average over all the subjects. In the

background, we show the A_z values from the discrimination analysis from each individual subject. Although any individual line is difficult to make out clearly, the purpose of this is to show the variation in A_z time-series across subjects.

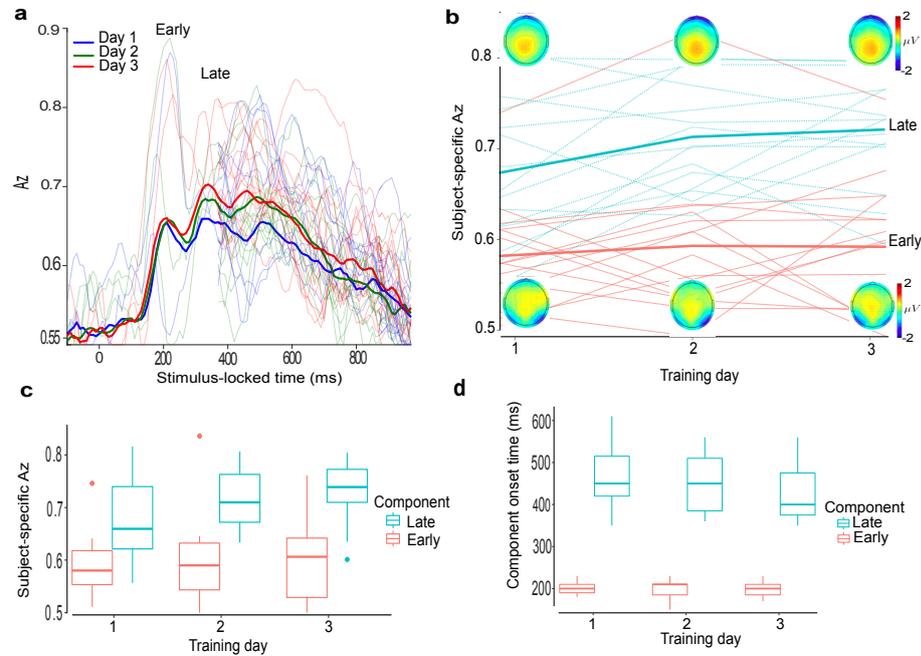


Figure 4.2: **Effects of prior information on perceptual learning and decision making.** (a) Multivariate discriminator performance (A_z) during face-versus-car outcome discrimination of stimulus-locked EEG responses across the three training days (1: blue; 2: green; 3: red), averaged across all 16 subjects, showing the presence of the early and late components. Faint lines represent individual subject data. (b) Average discriminator performance and scalp topographies for the early (blue) and late (light red) components across the three training days estimated at the time of subject-specific maximum discrimination (A_z). Faint lines represent individual subject data. (c) Boxplots showing the average discriminator performance for the early (blue) and late (light red) components across the three training days estimated at the time of subject-specific maximum discrimination (A_z). (d) Boxplots showing the average onset times for the early (blue) and late (light red) components across the three training days.

As can be seen, there are two¹ peaks or *temporal components* on each day (the bootstrap calculated threshold for significance at the $p < 0.01$ level is 0.55). These are iden-

¹It may be arguable that there is, in fact, a third component in this figure, e.g., around 500-600ms. This is largely an artifact of the averaging process over all the subjects on all the days because some components

tical to the components that we have described in details in Chapter 2 and in Diaz et al. (2017). The first is the early temporal component and this occurs around 200 ms. As described previously, this is related to the sensory encoding found in previous work (see, e.g., Philiastides & Sajda, 2006; Philiastides et al., 2006a; Ratcliff et al., 2009b; Philiastides et al., 2006a). The other temporal component is the late temporal component and this occurs around 400 ms. As described previously, this is related to post-sensory decision making (see, e.g., Philiastides & Sajda, 2006; Philiastides et al., 2006a; Ratcliff et al., 2009b; Philiastides et al., 2006a).

From this discrimination analysis, for each subject on each day, we extract the A_z value and onset time of both the early and the late peak. We will first analyse whether the A_z value changes with day of learning for both the early and late components. We then analyse how the onset times of the early and late components differ.

Change in component magnitude by day

In Figure 4.2c, we provide boxplots of the distributions of A_z values for the early and late components on each of the three days of the experiment. In Figure 4.2b, we provide the same information via line plots with the scalpmaps superimposed.

We statistically analyse the change the A_z value by performing, for the early and late component, a separate random effects linear regression with day of learning as the predictor variable and varying intercept random effect for each subject. Specifically, for both the early and the late components, we use the following model:

$$A_{z_i} = \beta_0 + \beta_1 x_i + \underbrace{\gamma_{[\text{subject}_i]}}_{\text{random effects}} + \epsilon_i,$$

where A_{z_i} is the A_z value for a particular observation, x_i is the value of the day for that observation, and subject_i is the subject for that observation. We assume that the γ values are normally distributed with mean 0 and standard deviation of σ_{subject} . We then assess whether there is a significant effect of day for the early and late components.

for some subjects on some days occur a bit later than others. We can observe stronger early discrimination in some subjects than in others. However, we do not observe that the early component moves earlier in time like we have observed for the late temporal component in both experiments. These are simply individual differences for the early temporal component.

The fixed effects coefficients for the late component analysis are:

	Estimate	Std. Error	df	t value	Pr(> t)
(Intercept)	0.66	0.02	40.37	30.02	<0.01
day	0.02	0.01	29	2.91	0.01

and for the early component analysis, the corresponding results are:

	Estimate	Std. Error	df	t value	Pr(> t)
(Intercept)	0.59	0.02	38.61	25.67	<0.01
day	0.01	0.01	29	0.61	0.54

From this, we observe that the A_z value for the late component rises by approximately 0.02 per day, and that this is a significant effect, $t(29) = 2.91$, $p = 0.01$. For the early component on the other hand, the A_z value rises by approximately 0.01 per day, but this is not a significant effect, $t(29) = 0.61$, $p = 0.54$.

Change in component onset time by day

In the Figure 4.2d, we provide boxplots of the distributions of the onset times of the early and late components on each of the three days of the experiment. To statistically analyse the change in onset times per component, we again use random effects linear regression models, and perform separate analyses for the early and the late components. This model has an identical form to the one used for the analysis of the A_z values. i.e. day of learning is the predictor variable and there is a varying intercept random effect for each subject. Specifically, for both the early and the late components, the model is the following:

$$y_i = \beta_0 + \beta_1 x_i + \underbrace{\gamma_{[\text{subject}_i]}}_{\text{random effects}} + \epsilon_i,$$

where y_i now represents the onset value for a particular observation, with x_i being the value of the day for that observation, and subject_i is the subject for that observation. Again, we assume that the γ values are normally distributed with mean 0 and standard deviation of σ_{subject} . We also again assess whether there is a significant effect of day for the early and late components.

The fixed effects coefficients for the late component analysis are:

	Estimate	Std. Error	df	t value	Pr(> t)
(Intercept)	491.10	20.59	33.20	23.85	<0.01
day	-22.67	6.67	29	-3.40	<0.01

and for the early component analysis, the corresponding results are:

	Estimate	Std. Error	df	t value	Pr(> t)
(Intercept)	199.30	5.87	42.47	33.94	<0.01
day	≈ 0.00	2.32	29	<0.01	1.00

From this, we observe that the onset time for the late component changes by approximately -22.70 per day, and that this is a significant effect, $t(29) = -3.40$, $p = < 0.01$. For the early component on the other hand, the onset time changes by almost exactly 0 per day, and clearly this is not a significant effect, $t(29) \approx 0.00$, $p = 1.00$.

Change in discrimination variable by cue

As described in detail in Chapter 2, in this discrimination analysis, we used a logistic regression based multivariate discriminant analysis to identify the EEG patterns that reliably distinguish between face and car stimuli. In particular, during each time window t in all the experimental trials for a given subject on a given day, our aim was to find a $K = 64$ dimensional vector \vec{w}_t such that

$$P(z_t = \text{face}) = \frac{1}{1 + e^{-y_t}}, \quad \text{where } y_t = \vec{w}_t' \vec{x}_t,$$

maximizes the probability of a correct classification of the stimulus. Note that here, \vec{x}_t is the $K = 64$ dimensional EEG signal at time t . The y_t is a *scalar* quantity that we refer to as the *discriminating variable* at time t . As can be seen, it is simply a weighted average of the entire EEG signal at time t that will maximally enhance those patterns in the EEG signal that reliably distinguish between the processing of face and car stimuli. In other words, y_t is 1d summary of all the EEG activity at time t that is relevant to discriminate face from car stimuli.

Here, we analyse how the difference of y variable on the face and car trials changes with the predictive value of the cue, and whether this differs between the early and late components. The difference between the y variable on the face and car trials, measured at the early and the late component, is effectively a measure of the strength of those components because the greater the difference between the y_t on the face and car trials, the better the EEG signal discriminates between face and car stimuli.

For each subject, on each day, we obtain the y variable value for face and for car stimuli at the early and late component time points, and then calculate the difference between them. We then assess how this difference changes with the cue's predictive value. The cue's predictive value is how well the cue predicts the stimulus on any given trial. For each, if on a particular trial, the cue is 70F/30C, and the stimulus on that trial is a face, then the cue's predictive value is 0.70. Likewise, if the stimulus on a particular trial is a car and the cue is 30F/70C, then again the cue's predictive value is 0.70.

We use a random effects linear regression to predict the average difference between y variable for faces and y variable for cars as a function of predictive strength of cue and day of learning and their interaction. We perform two separate analyses, one for the y variable at the early component time and one for y variable at the late time. Here, we provide the AIC values for these two separate analyses:

Early component model	Late component model
60.77	36.47

As can be seen, the AIC value for the model based on the late component is lower and hence is a better fit than the early component model. In other words, there is a stronger relationship between the predictive power of the cue, i.e. the closeness of the match between the cue and the stimulus, and the strength of the late component than between the predictive power of the cue and strength of the early component.

We use a similar approach to see how the average standard deviation of the y variable varies with the cue and how this differs between the early and late components. We analyse standard deviations for identical reasons to those illustrated in, for example, Figure 2.4b-c (page 44) and explained in Section 2.4.3 (page 43). In other words, any im-

improvements in learning can possibly be explained by decreasing the noise and variability in the neural decision variable.

Again, we use a random effects linear regression to predict how the average standard deviation of the y variable for both faces and cars varies as a function of predictive strength of cue and day of learning and their interaction. Again we perform two separate analyses, one based on y values at the early component time and one based on y at the late time. Here, we provide the AIC values for these two separate analyses:

Early component model	Late component model
-286.40	-334.90

As can be seen, as was the case for the differences between y_t values, the AIC value for the model based on the late component is lower. In other words, there is a stronger relationship between the predictive power of the cue and the variability of the late component than between the predictive power of the cue and variability of the early component.

Analysis of the discriminating component and behavioural responses

In the analysis we describe in this section, we focus specifically on how the early or late y variable relates to the subject's behavioural response. Our aim in this analysis is to determine how the y variable varies with the cue on each day. We then compare this change with the change in a probability of a face response as a function of the cue.

In the Figure 4.3a, we show the y variable (shown on the y -axis) at the early and late component times during the face trials and the car trials, and show how these vary as a function of the cue probability on those trials.

We then compare the y variable values to the probability of the subject making a response of 'face' on any given trial. The probability of a face response as a function of the cue value is plotted in Figure 4.3b. In comparing the y variable and the probability of a face response, we aim to discover if the y variable at either the early or the late component times is correlated with the subject's behavioural response.

We perform separate analyses for the early and the late components, In both cases, the analysis is based on a random effects linear regression model similar to what we used

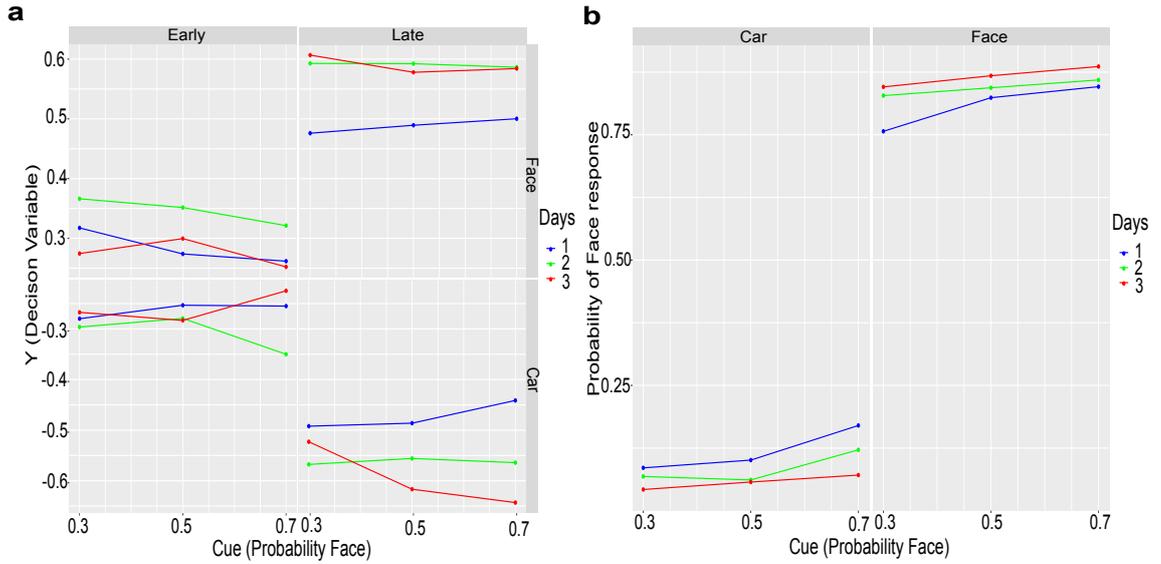


Figure 4.3: **Discrimination variable varies with pre-stimulus cue in perceptual decision making.** (a) Illustrates the change in y discriminating variable at the early and late component times as a function of the pre-stimulus cue, stimulus type (face; car), and days across 16 subjects. (b) The average probability that the subject will respond ‘face’ as function of the pre-stimulus cue, stimulus type, and days across 16 subjects.

above. Specifically, for this analysis, our model is

$$f_i = \beta_0 + \beta_1 y_i + \underbrace{\gamma_{[\text{subject}_i]}}_{\text{random effects}} + \varepsilon_i,$$

where f_i is the probability of ‘face’ response by a given subject on a given trial, y_i is the average value of that subject’s y variable at either the early or late component time on a trial on the same day, and having the same stimulus, and the same cue probability as the trial corresponding to f_i .

In order to compare how well the y variable at the early component time versus the y_i at the late component time predicts the probability of a face response, we evaluate the fit of both models. In both models, the outcome variable, i.e. f_i , is the same, so we can evaluate the predictive ability of the early and the late components easily by simply evaluating model fit. In both cases, we provide the AIC (defined above), the Bayesian Information criteria (BIC), and the Deviance. The AIC is defined and described on Page 73. BIC, on the other hand, is defined and described on Page 47. As mentioned, BIC is an approximation of the model’s log marginal likelihood and as such can be used to approximate the Bayes

factor comparing two models. We also report the Deviance. Deviance is simply -2 times the log of the likelihood function at its maximum likelihood value, i.e.

$$-2\log L.$$

The AIC, BIC, and Deviance of models are interpreted in relative terms, with the model with the lower value being the preferred model.

The AIC, BIC and Deviance of the two models are as follows:

	BIC	AIC	Deviance
early component	-97.84	-112.10	-120.10
late component	-169.20	-183.50	-191.50

In all cases, we see that the y variable value at the late component time is a markedly better predictor of the probability of a face response. These differences are all on the order of ≈ 70 . Given that either AIC or BIC differences of over 10 are considered substantial evidence in favour of the lower model, there is no question that the discriminating variable at the late component time is more highly related to the behavioural response.

Single-trial EEG discrimination analysis: Predictive versus non-informative cues

EEG single trial discrimination analysis may be applied not just to discriminating between the two type of stimuli, but between any two mutually exclusive classifications of all trials. In particular, we may perform a discrimination analysis to identify the EEG temporal components of the recognition of pre-stimulus cues. The aim of this analysis is to determine if there is a post-cue but pre-stimulus change in neural activity. This would imply a change in baseline activity prior to stimulus onset, as described by Basso and Wurtz (1998); de Lange et al. (2013); Albright (2012); Puri et al. (2009).

In order to do this, we divide all trials into two class: trials where there is a predictive pre-stimulus cue, trials where there is a non-informative cue. The 50F/50C cue that we have used in our analysis is clearly non-informative. In other words, it provides no predictive information about the upcoming stimulus. On the other hand, the 70F/30C and

30F/70C both provide usable information about the upcoming stimulus. Both tell us that there is a high probability of a particular stimulus occurring, and a low probability of the other stimulus.

Here, we perform a EEG discrimination analysis between the 50F/50C trials versus the 70F/30C and 30F/70C trials². Furthermore, we perform this discrimination analysis on the cue-onset locked trials. In other words, we perform the discrimination analysis over the period of the processing of the pre-stimulus cue rather than the processing of the stimulus itself. Other than the change the discrimination classes, and the definition of the trial onset time and duration, the procedure for this predictive versus non-informative discrimination analysis is the same. In other words, for each subject on each day and each time window t in the trial we aim to find a $K = 64$ dimensional vector w_t that maximizes

$$P(z_t = 70F/30C) = \frac{1}{1 + e^{-y_t}}, \quad \text{where } y_t = \vec{w}'_t x_t,$$

or the probability of correctly classifying each cue as either predictive or non-informative.

In the Figure 4.4a, we show the A_z values from onset of the pre-stimulus cue to 1400ms later, on each day of training. As above, we show the average over the subjects in the darker lines and the lighter lines show the A_z values for the individual subjects, which allows us to see the extent of inter-subject variability. What is clearly evident in this discrimination analysis is a single peak at around 200ms. Moreover, this peaks appears to increase in magnitude and occur later in time with each training day. In Figure 4.4b, we plot the average EEG activity, where we average over all trials and all subjects on each day, over the frontal electrode Fp2. This electrode was chosen as frontal site had highest activity in the scalpmap for the discriminator (as can be seen in the scalpmaps in Figure 4.4c. There is a clear negative (note, following convention, the y-axis is reversed such that negative values extend upwards) deflection in the EEG activity at this site at approximately the same time as the EEG discrimination component in Figure 4.4a.

In Figure 4.4c, we show the A_z values at the peak component on each day. Again, the dark line is the average, with the lighter lines being the individual subjects. We su-

²Note that while it is, in principle possible to also compare, for example, 70F/30C versus 30F/70C trials, in practice, this is problematic because of the insufficient number of trials of each class. More importantly, our immediate focus here is on informative versus non-informative cues, rather than cues with equally informative values.

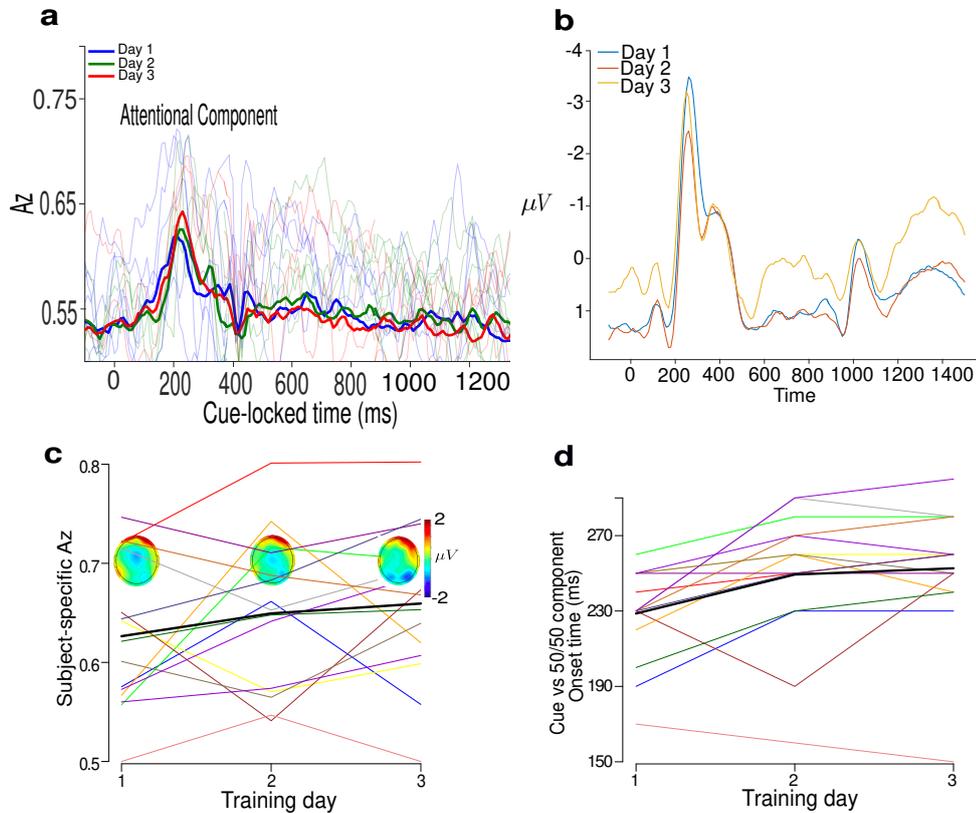


Figure 4.4: **Post-cue & pre-stimulus change in neural activity during perceptual decision making.** (a) Multivariate discriminator performance (A_z) during informative vs non-informative discrimination of stimulus-locked EEG responses across the three training days (1: blue; 2: green; 3: red), averaged across all 16 subjects. Faint lines represent individual subject data. (b) The cue-onset locked average EEG activity at frontal electrode site Fp2. This frontal site was chosen as frontal sites had highest activity corresponding to the discriminator. The EEG activity at this electrode was averaged over all trials and over all subjects on each day. (c) Average discriminator performance and scalp topographies for the peak of the component seen in (a), across the three training days. Colour lines represent the 16 individual subject data. (d) Average pre-stimuli (30/70;70/30) vs neutral (50/50) onset time for the component (black line) across the three learning days. Colour lines represent the 16 individual subject data.

perimpose the scalpmaps at the peak times on each day. These scalpmaps show higher activation in the attention areas.

We analyse whether there is a significant increase in the A_z value with training by using a linear random effects regression. Similarly to above, our model here is:

$$A_{z_i} = \beta_0 + \beta_1 x_i + \underbrace{\gamma_{[\text{subject}_i]}}_{\text{random effects}} + \varepsilon_i,$$

where x_i is the value of the training day, and $\gamma_{[\text{subject}_i]}$ is a random effect for the subject, with these random effects assumed to be normally distributed with zero mean and standard deviation σ_γ . The coefficients for the fixed effects are as follows:

	Estimate	Std. Error	df	t value	Pr(> t)
(Intercept)	0.61	0.02	39.40	25.11	<0.01
day	0.02	0.01	29	1.85	0.07

As can be seen, the effect of training on the magnitude of the peak is only marginally significant. We can confirm this by, in addition, performing a log-likelihood ratio test comparing the fit of this model to the null. This too is only marginally significant: $\chi^2(1) = 3.36$, $p = 0.07$.

In Figure 4.4d, we plot the peak onset times of the component per day. Here, the darker line is the average, and also shown are the lines for each subjects. To asses whether the peak occurs later in time with training, we again employ a random effects linear regression. It has an identical form to the regression model for the A_z value, with the only difference being that the outcome variable is now the component peak time rather than its A_z value. The coefficients of the fixed effects are as follows:

	Estimate	Std. Error	df	t value	Pr(> t)
(Intercept)	219.60	9.65	31.97	22.76	<0.01
as.numeric(day)	12	3.05	29	3.94	< 0.01

As can be seen, the effect of training on onset time is significant. We can also confirm this by performing a log-likelihood ratio test comparing the fit of this model to the null. This too is significant: $\chi^2(1) = 12.85$, $p < 0.01$.

4.4 Discussion

The main results from this study can be summarized as follows:

1. We have obtained strong evidence that both the speed and accuracy of perceptual decisions increased with the increasing predictive power of the pre-stimulus cue. This effect is observed at both noise levels, for both types of stimuli, and on all days of learning.
2. On the basis of the behavioural data, although speed and accuracy both improve significantly with each day of learning, the role of the predictive power of the pre-stimulus cue on perceptual decision making does *not* appear to change over the course of the three days of learning.
3. From the single-trial EEG analysis, we have replicated our main results described in Diaz et al. (2017) and Chapter 2. In particular, we have again identified two temporal components to perceptual decision making; an early and a late component. The strength of the late component, but not the early one, increases with each day of learning. Likewise, the late component, but not the early one, occurs at earlier onset times as learning proceeds.
4. The predictive power of the cue affects the average distance between y decision variable for faces and cars at the late component time more than at the early component time. Likewise, the predictive power of the cue affects the average standard deviation of y variable more during the late component than during the early component time.
5. The correlation between the y decision variable and the probability of a face response is greater during the late component time than during the early component time.
6. We have observed a temporal component at around 200ms after the onset of the cue that signifies whether the cue is informative or non-informative. There is a marginally significant increase in the strength of this component over the three days. It occurs significantly later in time as learning proceeds.

We will now discuss each one of these findings in turn.

The behavioural results clearly demonstrate that the prior expectations based on pre-stimulus cue affect perceptual decision making. As the probability of the upcoming stimulus according to the cue increases, so too does the speed and the accuracy of the perceptual decision. We observe this effect on speed and accuracy on every day of the experiment, and it equally affects the recognition of the face and the car stimuli, and at all noise levels. The importance of this result is that it confirms that the behavioural experiment worked as planned. The phenomenon that we expected to observe based on the work of, for example, Palmer (1975); Davenport and Potter (2004); Auckland et al. (2007); Saffran et al. (1996); Aslin et al. (1998); Fiser and Aslin (2002); Bar (2004); Oliva and Torralba (2007), and others, was indeed observed. As such, we may interpret the EEG and other results with more confidence knowing that the experimental paradigm was methodologically sound.

One important unexpected result that emerged from the behavioural analysis was that the effect of the pre-stimulus cue did not change over the course of the three days. This emerged clearly from the regression analysis of both the speed and the accuracy data. In both of those analyses, the effect of cue did not interact with the effect of training day. There were clear and strong *main* effects of both cue and day, but, for both the reaction time and the accuracy analysis, the models with the interaction between these two predictors had higher AIC values than the corresponding models without the interactions. This was an unexpected result as we predicted that there would be a *decrease* in the effect of the cue with increased training. As described above, this was predicted from the Bayesian principle that prior information should influence decision making more when sensory evidence is weaker (e.g. Bogacz et al., 2006). According to our explanation of the results presented in Chapter 2, which equally hold here, training leads to a faster and more efficient accumulation of information. As such, sensory evidence will be stronger by, for example, the third day of training compared to the first. Hence, the effect of the prior should be relatively weaker on Day 3 than Day 1. However this effect was not observed. Moreover, we did not observe any interaction, either in the accuracy or the reaction time analyses, between the cue and the level of noise of the stimulus. Again, such an interaction is directly predicted by the Bayesian account of perceptual decision making whereby the prior probability of a stimulus becomes less informative the stronger

the sensory evidence in favour of the stimulus. While this results may be taken as a falsification of the Bayesian model, it is also arguable, however, that these effects were not observed simply due to low statistical power of the experiment. There were only 16 subjects in this experiment, and hence only relatively large statistical effects would have power in this situation. Interaction effects will be, by necessity, more subtle than main effects and hence would require larger sample sizes to observe.

The single-trial EEG analysis replicated the main findings that we reported in Chapter 2 and Diaz et al. (2017). In particular, the early and late temporal components were observed here as before. Likewise, we also observed again how the late component, but not the early one, became stronger and occurred earlier in time with increasing training over the three days. This replication is obviously a reassuring result and is a reflection on the robustness of the main results were reported in Chapter 2 and Diaz et al. (2017). This is especially the case given that there was an additional variable being manipulated in this experiment.

As explained previously, The y discriminating variable can be viewed as the EEG measured accumulation to bound decision variable. In other words, it plays a role similar to the stochastic variable in a sequential sampling model, such as a DDM, that moves towards one or another threshold on the basis of accumulating sensory evidence. As we described in Chapter 2 for discrimination between two perceptual classes to be effective, the average of the y variable when it represents each of the two classes should be further apart, and there should be less variability in the y values. We expect, therefore, that if the late component and not the early component is being affected the prestimulus cue, we ought to see the distance between the y values for faces and cars becoming further apart, and the standard deviation of these values becoming narrower, during the late component time period and not during the early component. There was a stronger prediction of the difference between the y variable values for the two classes on the basis of the cue during the late component than during the early one. Likewise, there was a stronger prediction of the standard deviation change with the cue during the late component time than the early component time. In a related result, we showed that how the cue affects the y variable parallels how the cue affects the probability of making a face response, but that this only occurs at the late component time and not the early component time.

The significance of these two results is that they point to the effect of the cue on perceptual decision as taking places at the late component time and not the earlier one. This result is consistent with the cue affecting the accumulation of the decision to reach a threshold, rather than affecting the pre-stimulus baseline activation. In other words, it is consistent with how prior expectation affects evidence accumulation, as revealed by, for example, Hanks et al. (2011); Cravo et al. (2013); Dunovan et al. (2014), and is less consistent with the baseline activation hypothesis as forwarded by, for example, Basso and Wurtz (1998); de Lange et al. (2013); Albright (2012).

Finally, we have observed a EEG component that discriminates between informative and non-informative cues that occurs roughly 220ms after the onset of the pre-stimulus cue (see Figure 4.4a, with the corresponding average EEG activity in Figure 4.4b). This component may be related to increased attentional focus that occurs when the cue is informative as opposed to when it is non-informative. Indeed, this temporal component may also be related to a component described in Philiastides et al. (2006a). This component arose in a single-trial analysis that discriminated between easy and difficult trials, and occurred at approximately 220ms after stimulus onset. It was argued in Philiastides et al. (2006a) to be an attentional component. In our case, our post-feedback component also occurs at 220ms after the cue onset. This is further indication therefore that this may be an attention related component.

A final point to discuss is how pre-stimulus cues compare with correct/incorrect feedback on trials. In the experiments described in this chapter, subjects performed a very similar task to what was described in Chapter 2. The only differences between the two sets of experiments is that in the experiments described here, subjects obtained a pre-stimulus cue that indicated (probabilistically) the upcoming stimulus identity, while in the experiments in Chapter 2, subjects obtained post-stimulus feedback concerning their accuracy. In both studies, we observe clear behavioural evidence for learning. In both cases, this is manifested both in terms of accuracy and reaction time. Moreover, across the two studies, the learning effects are roughly comparable. We may take this as evidence that feedback plays as much a role in learning as does the pre-stimulus cue.

Chapter 5

Hierarchical diffusion model analysis of perceptual decision making

5.1 Drift diffusion models

The DDM is a mathematical model of decision making under time constraints. To understand how the DDM works, it is helpful to start with a very simple mathematical model, namely the random walk, which is in fact the basis of the DDM.

Consider a variable $x \in \mathbb{R}$ that takes on different integer values at discrete time steps, i.e. at times $t = 0, t = 1, t = 2, \dots$. At time $t = 0$, we will set x to some integer value z . We'll denote the value of x at any time t by x_t , and so x at time $t = 0$ is denoted x_0 . After that and then for all $t \in \{1, 2, \dots\}$

$$x_{t+1} = \begin{cases} x_t + 1, & \text{with probability } p, \\ x_t - 1, & \text{with probability } 1 - p. \end{cases}$$

In other words, x starts with the value of z and then moves either 1 higher or 1 lower along the number line at each tick of some clock. We can now place two *barriers* at integer values on the number line. We'll denote these by a_l and a_u , with a_l being lower than the starting value z and a being above this starting value, i.e. $a_l < z < a - u$. Whenever the x variables reaches one of these barriers, it stops and does not change its value. In other

words, a_l and a_u are *absorbing* barriers. If x reaches a_u first, another random variable y takes on the value of 1, while if x reaches a_l first, y takes the value of 0.

On the basis of what we've just described, we can describe the process of the random walk. The variable x always starts at z and then randomly moves up or down in steps of 1. It moves up with probability p and moves down with probability $1 - p$. Continuing like this, eventually it will reach one of the two barriers, and then the process terminates. Clearly, because taking a step up or down is random, it could always reach the upper or the lower barrier first. Likewise, it could reach either barrier relatively quickly or wander up and down for some time before he reaches one or the other. See Figure 5.1 for a further illustration of this random walk process. Note that there are a set of *parameters* that control the random walk process, i.e., p, z, a_l, a_u . If these are set to constant values, then there will always be a constant probability of first crossing the upper or the lower barrier, and also there will always be a fixed probability distribution over how long it takes to reach each one barrier or another.

We can now extend this process as follows. First, and without loss of generality, we can always set a_l to be 0, and then rename a_u to a , and so the starting position is always $0 < z < a$. Then, instead of taking a step of one unit at each tick of the clock, we instead take a step size of

$$\Delta \triangleq \sqrt{\tau},$$

at every τ time interval. We also define p as

$$p \triangleq \frac{1}{2} (1 + \mu\sqrt{\tau}),$$

and so

$$1 - p \triangleq \frac{1}{2} (1 - \mu\sqrt{\tau}).$$

Now, in the limit of $\tau \rightarrow 0$, the above random walk process converges to a DDM with a drift rate of μ , an upper barrier of a , and initial starting point of z . In other words, the DDM is a continuous time version of the random walk. To be more precise, the DDM is random (or stochastic) process on a 1d space and in continuous time, i.e. the value of the variable changes smoothly or continuously in time, and does not just take discrete steps at discrete time points (or ticks). The rate of change of the random variable in the DDM is

$$\frac{dx_t}{dt} \sim N(\mu, 1),$$

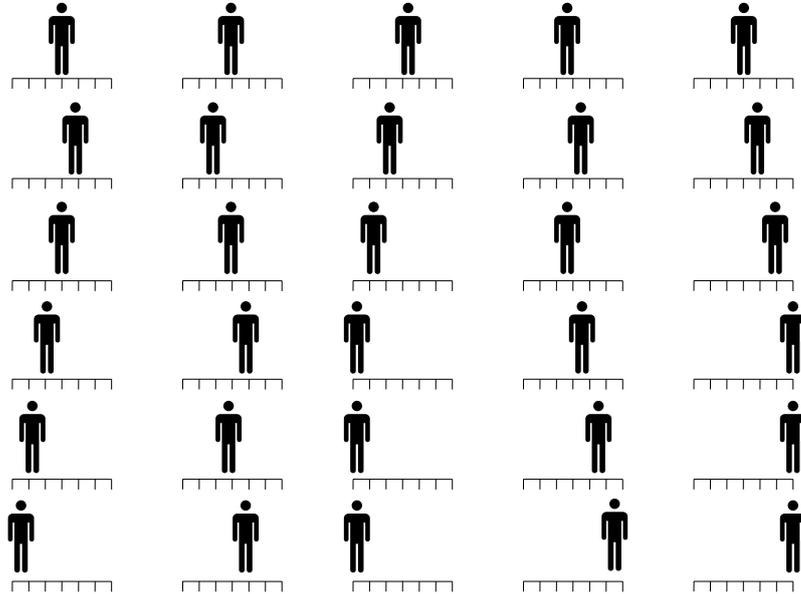


Figure 5.1: **Diffusion model as random walk** In this figure, we illustrate a simple concept underlying the DDM, namely the random walk. Consider a man standing at the center of a line on the floor. At each tick of the clock, he takes a step randomly to the left or to the right. He does this by, at each time step, flipping an unfair coin whose bias to come up heads is exactly p . If the coin comes up Heads, he steps left; if it comes up Tails, he steps right. As such, the probability that he steps left at any time is always exactly p , and the probability that he steps right is always exactly $1 - p$. Starting in the centre of the line, flipping the coin at each time step, and then stepping left or right each time, he eventually reaches the left or the right end of the line. Once he does so, he stops and remains there forever. Shown here are six repetitions of this random walk. For example, on the first repetition, shown in the leftmost column of figures, the man starts in the centre, steps to his left, then to his right, and again then to his right again, again, again. At this point, he has reached the rightmost edge of the line. On the second repetition, shown in the second column, he starts in the center and then steps to his right, then left, left, right, left, and does not reach either side in the depicted steps. In the third repetition, he reaches the right side by the third step. And so on.

where $N(\mu, 1)$ signifies a normal distribution with a mean of μ and variance of 1¹.

According to the DDM with parameters a , μ , and z , the probability of first reaching the upper barrier is

$$P(y = 1) = \frac{\exp(-2z\mu) - 1}{\exp(-2a\mu) - 1}$$

The probability density function over the time to reach the upper barrier is

$$P(t|y = 1) = f(t) = \frac{\pi}{a^2} \exp\left((a-z)\mu - \frac{\mu^2}{2}t\right) \frac{1}{P(y = 1)} \sum_{m=1}^{\infty} \sin\left(\frac{\pi m(a-z)}{a}\right) \exp\left(-\frac{1}{2} \frac{\pi^2 m^2}{a^2} t\right),$$

and the probability to reach the lower barrier is identical to the above function if we substitute μ for $-\mu$ and z for $a - z$.

In the DDM, it is common to also add a fourth parameter known as the *non-decision time*, which we will denote here by t_{Δ} . The effect of this parameter is simply to shift the probability density function $f(t)$ forwards in time. In other words, with the non-decision time parameter included we now have

$$P(t|y = 1) = \begin{cases} 0, & \text{if } t < t_{\Delta}, \\ f(t - t_{\Delta}), & \text{if } t \geq t_{\Delta}. \end{cases}$$

A figure illustrating the role of the parameters is shown in Figure 5.2.

5.1.1 Drift diffusion models of decision making

The above mathematical model can be used as a model of speeded decision making in a two-alternative forced choice task. According to this perspective, the choice that is made in this task and the response time taken to make this choice can be modelled by the probability of crossing the upper or lower barrier in a DDM. The random diffusion of the decision variable x represents the *accumulation of evidence* for one or the other of the two alternative choices. Different settings of the parameters make one choice or another and the response times for these choices are more or less likely. We will now explain these in turn.

¹Setting the variance of the Normal distribution to be 1, rather than allowing it to be a variable σ^2 that takes on different values does not limit the generality of the DDM. The variance of the drift rate is coupled to the value of a and μ , so that an identical DDM will result for any value of the variance if the value of a and μ are adjusted. As such, we can set the variance to a constant value, and it is convenient to set it to exactly 1 to simplify the probability distributions for time to reach the barriers, etc.

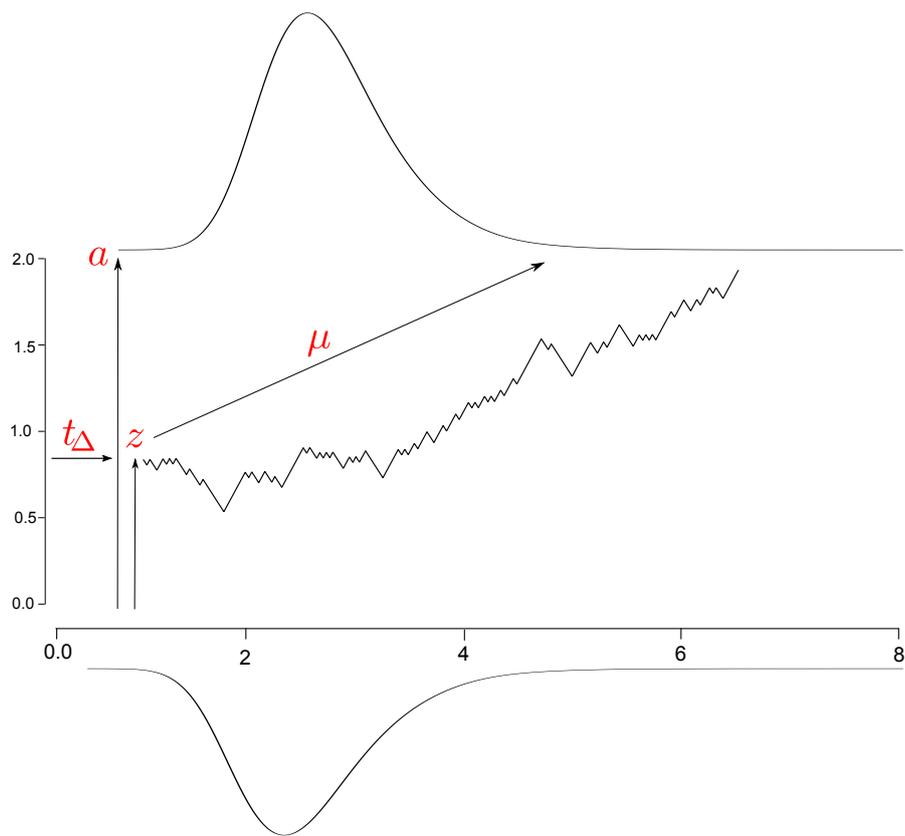


Figure 5.2: **An illustration of the DDM parameters.** Here, the random walk trajectory is an example of single diffusion process from a starting location z . The distance between the two barriers is a . The average drift rate is given by μ . The non-decision time is given by t_{Δ} . At the top and bottom are two probability density functions showing the probability of crossing the upper and lower barriers first, respectively.

- Drift rate, i.e. μ : If there is, on average, more evidence in favour of choice $y = 1$, then the drift rate μ will be positive and the diffusion process will be more likely to evolve over time in the positive direction, and so hit the upper barrier at a first. If there is more evidence in favour of the choice $y = 0$, then the drift rate μ will be negative and the diffusion process will be more likely to evolve over time in the negative direction, and so hit the lower barrier at 0 first. In general, the greater the absolute value of the drift rate, i.e. $|\mu|$, the faster the decision variable will reach a barrier. As such, assuming constant values of the remaining parameters, the drift rate μ controls the probability of reaching one barrier or the other first and controls

the distribution over the how long it takes to reach each barrier. Put in terms of decision making, this entails that the drift rate controls the probability of making a response $y = 1$ or $y = 0$, and on average how quickly these decisions are made.

- Increasing the value of a trades off speed for accuracy. The larger the value of a , assuming z begins equidistant between 0 and a , then if $\mu > 0$, the decision variable is always more likely to hit a before it hits 0, and vice versa if $\mu < 0$. With increasing a , the time taken to reach either barrier will increase, but it is less likely that the decision variable will cross 0 first if $\mu > 0$, or cross a first if $\mu < 0$. As such, a is the speed-accuracy trade-off parameter.
- The starting position z controls how close the diffusion model is to each one of the two barriers when it begins. As such, it represents a baseline bias in favour of the response represented by the barrier it is closest to. Put another way, the starting position represents the *a priori* evidence in favour of one response over the other before any sensory information is observed. For this reason, z is often also called the *bias* of the DDM.
- The non-decision time t_{Δ} represents the time taken for non-decision related processes, particularly sensory encoding time and motor response time. In other words, before evidence for one response or the other can even begin to accumulate, the sensory information must be encoded. Likewise, even after a decision has been made, it takes a non-zero amount of time to execute a motor response such as pressing a key on the keyboard. These processes are assumed to be constant values in any decision and so are represented simply by a shift to the right in the probability density function over response time.

5.1.2 Modelling behavioural data with drift diffusion models

Let us imagine that we have a choice and reaction time data from a behavioural experiment. For example, we could have n choices and reaction times all from one participant in one experiment where the participant's task is to determine if a noisy image is a face or

a car. The behavioural data that we have is

$$y_1, y_2 \dots y_i \dots y_n, \quad t_1, t_2 \dots t_i \dots t_n,$$

where y_i is the choice at trial i and t_i is the reaction time at trial i . We also have predictor variables

$$s_1, s_2 \dots s_i \dots s_n, \quad q_1, q_2 \dots q_i \dots q_n,$$

where s_i codes the identity of the stimulus at trial i , i.e., whether the stimulus is actually a face or car, and q_i represents the level of noise in the stimulus on trial i .

Recall that the common approach to statistical analysis of data like would be use to general or generalized linear models. For example, it would be common to model the response data using binary logistic regression, such as

$$\log \left(\frac{P(y_i = 1)}{P(y_i = 0)} \right) = \alpha + \beta_s s_i + \beta_q q_i, \quad \text{for all } i \in \{1, 2 \dots n\},$$

or to model the reaction times using general linear models as follows,

$$t_i = \alpha + \beta_s s_i + \beta_q q_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad \text{for all } i \in \{1, 2 \dots n\}.$$

(Note that in the above two equations, we used α , β_s , and β_q in both. This is not to say that we assume that the coefficients in both models are identical. We are simply (over-)using the same labels to avoid a proliferation of symbols.)

Statistical modelling of response data with a DDM is essentially an extension of this process. In particular, we model

$$y_i, t_i \sim \text{ddm}(\mu_i, z_i, a_i, t_{\Delta i}), \quad \text{for all } i \in \{1, 2 \dots n\},$$

where $\text{ddm}(\mu_i, z_i, a_i, t_{\Delta i})$ is a DDM with drift rate μ_i , starting position z_i , upper barrier position (and so inter-barrier distance) a_i and non-decision time $t_{\Delta i}$. In other words, we model the choice y_i and the reaction time t_i at any given time as a sample from a DDM with parameters $\mu_i, z_i, a_i, t_{\Delta i}$.

We can now model the parameters on each trial as a function of the predictors. For example, we could model μ_i using a general linear model such as

$$\mu_i = \alpha + \beta_s s_i + \beta_q q_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad \text{for all } i \in \{1, 2 \dots n\}.$$

What we are saying by this is that, for each trial i , the choice y_i and the reaction time t_i is a sample from a DDM with parameters $\mu_i, z_i, a_i, t_{\Delta i}$, and where each μ_i is a general linear model of the two predictors s_i and q_i .

In the above example, for simplicity we assumed that the starting position, inter-barrier distance, and non-decision times were known constants. Of course, in practice, we rarely if ever know the values of these parameters. Also, these values need not be constants across all trial. Thus, in general, we would provide probabilistic models for these parameters, and how they vary from trial to trial, too. To do so, it is first convenient to re-parameterize the starting position z as a proportion of the inter-barrier distance, i.e. instead of dealing with the parameter z , we instead deal with $b = \frac{z}{a}$. Note that with this transformation, $b \in (0, 1)$. With this change, our new model is

$$y_i, t_i \sim \text{ddm}(\mu_i, b_i, a_i, t_{\Delta i}), \quad \text{for all } i \in \{1, 2 \dots n\},$$

with, as before,

$$\mu_i = \alpha + \beta_s s_i + \beta_q q_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad \text{for all } i \in \{1, 2 \dots n\}.$$

Now we also provide probabilistic models for b , a , and t_{Δ} . For example, we would use a logistic regression model for the value of b_i . For example,

$$b_i = \alpha + \beta_s s_i + \beta_q q_i, \quad \text{for all } i \in \{1, 2 \dots n\}.$$

For the inter-barrier distance a_i , a suitable model could be

$$a_i \sim N_{[0, \infty)}(\alpha + \beta_s s_i + \beta_q q_i, \sigma_i^2), \quad \text{for all } i \in \{1, 2 \dots n\},$$

where $N_{[0, \infty)}$ signifies a *half-normal* distribution from 0 to ∞ , i.e., an other Normal distribution but where the probability of a value being less than 0 is 0, and so with probability 1, all values are non-negative. Similarly, we could model the non-decision time by

$$t_{\Delta i} \sim N_{[0, \infty)}(\alpha + \beta_s s_i + \beta_q q_i, \sigma_i^2), \quad \text{for all } i \in \{1, 2 \dots n\}.$$

(Note that we are again over-using the symbols α , β_s , β_q , σ^2 for the coefficients in all of the above models simply to avoid a proliferation of symbols.)

5.1.3 Hierarchical diffusion models

Just as we employed hierarchical or multilevel or random-effects regression models in Chapters 2, 3, and 4, so too can we extend the DDM to be a HDDM². This allows us to model inter-participant or inter-stimulus variability in the effect of the predictors on the observed data. For example, in the above model, we modelled the drift rate by

$$\mu_i = \alpha + \beta_s s_i + \beta_q q_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad \text{for all } i \in \{1, 2 \dots n\}.$$

However, this was under the assumption that the participant on all of the n trials was the same. Normally, of course, we will have multiple subjects and we do not usually wish to assume that the effect of the predictors on the choice and reaction time variables is identical across different subjects. To allow for variability, the multilevel approach could assume

$$\mu_i = \alpha_{\text{subject}_i} + \beta_{s,\text{subject}_i} s_i + \beta_{q,\text{subject}_i} q_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad \text{for all } i \in \{1, 2 \dots n\},$$

where

$$\alpha_j \sim N(\alpha_0, \tau_\alpha^2), \quad \beta_{s,j} \sim N(\beta_{s,0}, \tau_{\beta_s}^2), \quad \beta_{q,j} \sim N(\beta_{q,0}, \tau_{\beta_q}^2) \quad \text{for } j \in 1 \dots J$$

In other words, we are assuming subject-specific coefficients, for all J subjects, that are themselves drawn from population general probabilistic models. Using the sample principles, we can just as easily apply the hierarchical modelling approach to the models for b_i , a_i , and $t_{\Delta i}$.

5.1.4 Bayesian inference using Gibbs sampling

For all the HDDM statistical analysis we employ here, we implement our models in the Just Another Gibbs Sampler (JAGS) (*Just Another Gibbs Sampler*) developed by Martin Plummer and others, and available under GNU Public Licence at

<http://mcmc-jags.sourceforge.net>.

²This type of analysis is a relatively recent innovation. Methodological papers that describe this analysis and provide software implementations include, for example, Vandekerckhove, Tuerlinckx, and Lee (2011); Wiecki, Sofer, and Frank (2013), while empirical studies in the general field of cognitive neuroscience that have used the HDDM for analyses include Zhang and Rowe (2014); Dunovan et al. (2014).

JAGS is a probabilistic modelling language that automatically derives a Gibbs sampler Markov Chain Monte Carlo (MCMC) that allows samples to be drawn from the posterior distribution over all unknown variables, i.e. parameters, latent variables, etc., in the probabilistic model, and this allows us to infer the probable values of these unknowns from the data. In other words, JAGS allows us to perform statistical inference of parameters etc. for arbitrary probabilistic models. The definition of the probability distribution in the DDM that are implemented in JAGS were provided as extension modules by Joachim Vandekerckhove and others, and also available under GNU Public Licence at

<https://sourceforge.net/projects/jags-wiener>.

5.1.5 Advantages of diffusion models

The standard approach to the analysis of reaction time and choice behavioural data is either identical to or similar to the methods we have used for analysing reaction time and accuracy in, for example, Chapter 2 and Chapter 4. There, we analyse reaction times using (multilevel) linear regression models and separately we analyse accuracies using (multilevel) logistic regression. Diffusion models, and sequential sampling models generally, offer advantages over these methods. For example, viewed solely as methods for statistical analysis of reaction time and choice data, diffusion models make more realistic assumptions about the nature of reaction times, and the relationship between reaction times and choice or accuracies. For example, traditional methods often assume reaction times are normally distributed around a mean that varies as a function of predictor variables. On the other hand, a diffusion model assumes that these reaction time distributions are unimodal with a positive skew and long positive tail. It is well demonstrated that human reaction times have these characteristics (see, e.g., Ratcliff & Smith, 2004; Ratcliff et al., 2004; Ratcliff & McKoon, 2008; Ratcliff & Van Dongen, 2011, for detailed reviews). Likewise, while separate analyses of reaction times and accuracies assume that these two outcome variables are independent of one another, it is well appreciated that these are coupled and can trade-off one another (see, for example Heitz, 2014, for a review of this topic). Diffusion models treat reaction and choice or accuracy as two coupled outcome variables whose distributions are determined by a common set of parameters, and where

speed-accuracy trade-off in particular is determined by one of these parameters (that is, inter-barrier distance). In addition, diffusion models also have mechanistic interpretations in terms of neural processes. We have described the link between sequential sampling models generally and the cognitive neuroscience of decision making in Section 1.0.2. For example, with respect to our particular studies, properties of diffusion models, such as the drift rate, have been shown by Philiastides et al. (2006b); Ratcliff, Philiastides, and Sajda (2009a) to relate the late EEG temporal component that we have described throughout this thesis (see Section 5.4 for more detail).

5.2 Diffusion model analyses of the effect of learning on perceptual decision making

In the present analysis, we model the response speed and choice data from our first project, i.e. the project described in Chapter 2 and in Diaz et al. (2017), using a HDDM. Our particular aim is to determine how learning affects perceptual decision making by examining which parameters of the diffusion model vary with training.

In this model, our outcome variables are the participant's choice, i.e. face or car, on each trial, and their reaction time to make the choice. These are modelled by first passage times in the DDM: Crossing the upper barrier signifies the choice of face, and the lower barrier signifies the choice of car. In the DDM, in general, the drift rate, inter-barrier distance, bias, and the non-decision time can all, in principle, vary by day of learning and by coherence on each trial. The model is a HDDM given that drift rate, inter-barrier distance, bias, and the non-decision time all vary randomly by subject.

5.2.1 Model selection

We begin by considering a range of possible HDDM models. These differ from one another in terms of which of the main predictor variables affect which of the main parameters of the HDDM. In all models, we assume that the starting position is a fixed and known constant. Specifically, we assume that the $b = 0.5$. This assumption is very commonly made in drift-diffusion modelling (see, e.g., Farrell & Lewandowsky, 2018, Chapter 14),

unless there is a definite and strong possibility that the starting point will systematically vary with one of more predictor. In our case, because *face* and *car* are exactly equally likely in the experiment across all days and across all participants, and also these stimuli are generated using identical images, we should expect a symmetry in response to the two classes, and so the starting point of $b = 0.5$ is a reasonable assumption.

We then consider the following 7 variants of HDDM. However, in all cases, we assume the boundary separation values, the drift rate (for both face and car stimuli), and the non-decision time varies randomly across subjects, and we model these hierarchically.

1. Drift rate (for both faces and cars) varies by the training day, and by the coherence level. Boundary separation, and non-decision time are constant across these predictors.
2. Drift rate varies by the training day, and by the coherence level. Boundary separation varies by day, but is constant for coherence. Non-decision time is constant by both learning day and coherence level.
3. Drift rate varies by the training day, and by the coherence level. Boundary separation varies by coherence, but is constant for day. Non-decision time is constant by both learning day and coherence level.
4. Drift rate varies by the training day, and by the coherence level. Boundary separation varies by coherence, and by day. Non-decision time is constant by both learning day and coherence level.
5. Drift rate varies by the training day, and by the coherence level. Non-decision time varies by day, but is constant for coherence. Boundary separation is constant by both learning day and coherence level.
6. Drift rate varies by the training day, and by the coherence level. Non-decision time varies by coherence, but is constant for day. Boundary separation is constant by both learning day and coherence level.
7. Drift rate varies by the training day, and by the coherence level. Non-decision time varies by coherence, and by day. Boundary separation is constant by both learning day and coherence level.

Note that in all cases, we assume that drift rate varies by day and by coherence. In other words, we take it as given that the drift rate will vary by these two predictors, based on the noticeable change in both accuracy and reaction time as a function of these two predictors. Our main question, therefore, is whether boundary separation and non-decision time also vary by one or other of the two main predictors.

For each model, we calculate the Watanabe Akaike Information Criterion (WAIC) (see S. Watanabe, 2010) and use this for model evaluation. WAIC is calculated as follows:

$$\sum_{i=1}^n \log \left(\frac{1}{S} \sum_{s=1}^S P(y_i | \theta^s) \right) - \sum_{i=1}^n V_{s=1}^S (\log P(y_i | \theta^s)),$$

where y_i is the observed data at observation i , θ^s is a single sample from the posterior over all parameters (i.e. θ generically denotes all parameters in the model). Here, we assume n observations, and S samples from the posterior distribution over θ . The term $V_{s=1}^S(\cdot)$ signifies the variance of its arguments. WAIC is related to AIC in that it aims to provide a close approximation to leave one out cross validation, without having to perform the resource intensive repetitions of the model with different data-sets. However, it has been shown by S. Watanabe (2010) that WAIC is a more *widely* applicable approximation to leave one out cross validation. Moreover, WAIC is computed from MCMC samples and does not require calculation of the maximum of the likelihood function. Note that by choosing WAIC (or AIC or cross-validation) as the criterion for model selection, we are defining the *best* model as the model with the best out-of-sample predictive error.

For each of the seven models defined above, we infer the posterior distribution over all unknown variables using JAGS, and then from the samples, we calculate the WAIC. The results are

WAIC	model description
3271	drift \sim day + coherence
3510	drift \sim day + coherence; boundary \sim day
3309	drift \sim day + coherence; boundary \sim coherence
3510	drift \sim day + coherence; boundary \sim day + coherence
3642	drift \sim day + coherence; non-decision time \sim day
3326	drift \sim day + coherence; non-decision time \sim coherence
3637	drift \sim day + coherence; non-decision time \sim day + coherence

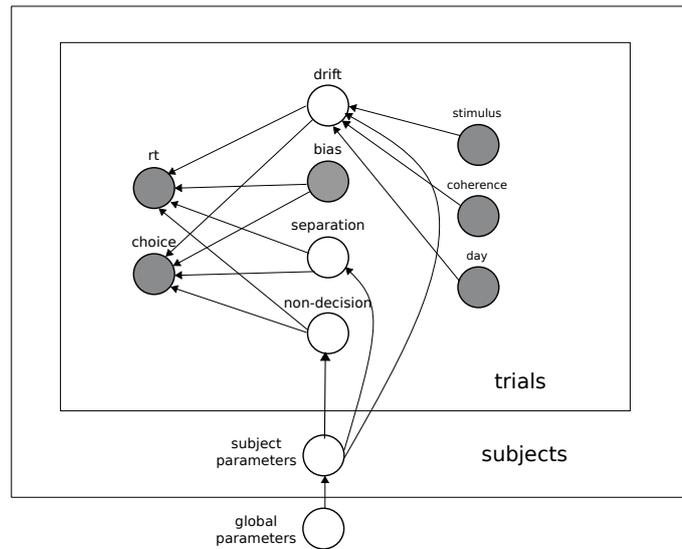


Figure 5.3: **A graphical model illustration of the best fitting model of the effect of learning on perceptual decision making.** Illustration of the best fitting model for the data from the perceptual learning experiment (see Diaz et al. (2017) or Chapter 2). This shows which predictors and other observed variables affect or influence which parameters in the system, and how these affect the accuracy and reaction time. Note that shaded nodes indicate observed data. Here, the bias variable is shaded because we have modelled it as a fixed and known constant.

As can be seen, the model with the lowest WAIC, and so therefore the model we choose, is the one where only the drift rate changes as a function of the main predictors, and the boundary separation and non-decision time are constants. Recall, however, that the effect drift rate, boundary separation, and non-decision time are all modelled as randomly varying across subjects. In what follows, we provide further results for this best fitting model alone. In Figure 5.3, we provide a graphical model illustration of this model.

5.2.2 Analysis of best fitting model

We begin by comparing the predictions of the model to the empirical data from the subjects. In Figure 5.4a and Figure 5.4b, we provide scatter-plots comparing model predictions to the subjects' reaction times, for both stimuli classes, coherence levels, and days of training, and for correct and incorrect reaction times. For the correct responses, the correlation coefficient, using the *robust correlation* method named `cov.rob` in the MASS R package, for the fits to the face and car stimuli are as follows:

Car	Face
0.74	0.80

For the *incorrect* responses, the correlation coefficient for the fits to the face and car stimuli are as follows:

Car	Face
0.62	0.74

In both cases, we can see that the models' reaction time predictions always resemble the empirical reaction for the different stimuli and on the different days. There are noticeably closer fits, however, for the face stimuli compared to car stimuli.

In Figure 5.4c and Figure 5.4d, we show the 95% HPD intervals for the boundary separation parameter and non-decision times, respectively, for each subject. In general, the 95% HPD interval for a variable gives the range of values between which lies 95% of the probability mass in a posterior probability distribution. In other words, it can be interpreted like a confidence interval, specifically telling us that there is a 95% probability that the true value of unknown variables lies in the defined interval. For each variable, the set of HPD values show the extent of the inter-subject variability in that variable. In particular, for the boundary separation variable, the HPD intervals show how subjects differ from one another in terms of their speed accuracy trade-off. For the non-decision time, the variability shows how subjects differ considerably in their sensory-coding and motor response times. In general, the fact that some subjects' HPD interval are considerably longer than others indicates the uncertainty we have about those subjects' values for this variable.

In Figure 5.4e, we show the posterior distribution over the drift rate for face and car stimuli on the different days of learning and for the different coherence levels. As can be seen, there is noticeable increase in the absolute value of the drift rate with each day of learning. In other words, for faces the drift rate increases with training, and for cars, the drift rate decreases. Recall, that a face response is modelled by crossing the upper barrier in a DDM, while the car response is modelled by crossing the lower barrier. As such, all face stimuli should correspond to positive drift rates and all car stimuli should have

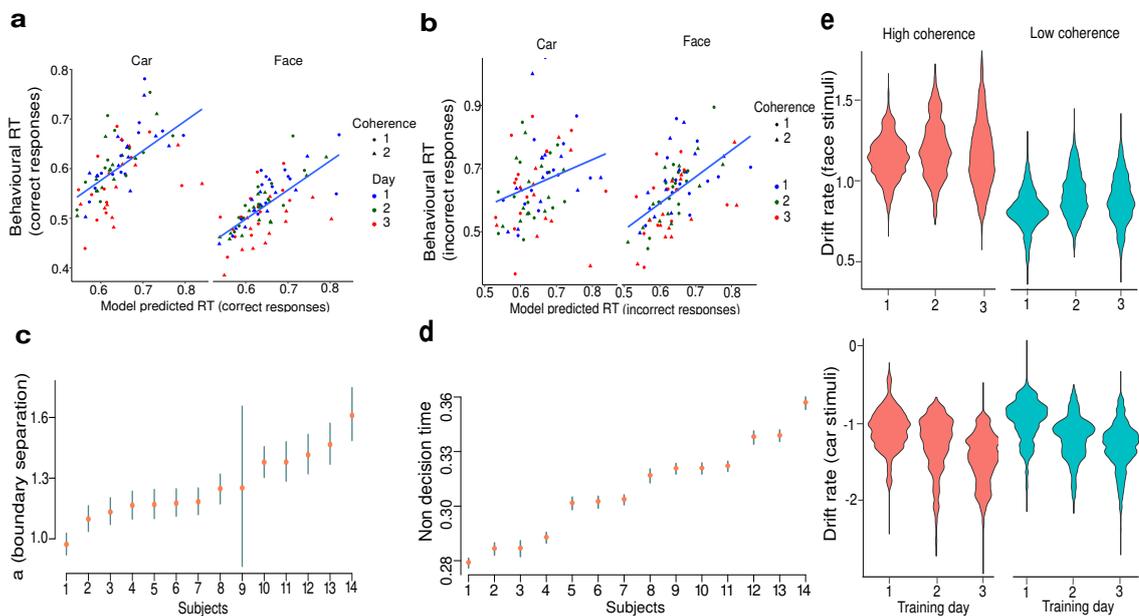


Figure 5.4: **Diffusion Model of the effect of learning on perceptual decision making.** (a) Illustration of model fit. The model’s predicted reaction times for stimuli (face; car), on each day and for the different coherence levels are compared to the average reaction times of all subjects for these variables. Here, we show reaction times according to the model and from the subject, where the responses was correct. (b) Illustration of model fit. The model’s predicted reaction times for stimuli (face; car), on each days and for the different coherence levels are compared to the average reaction times of all subjects for these variables. Here, we show reaction times according to the model and from the subject, where the response was *incorrect*. (c) Shows the 95% HPD interval of the Boundary Separation (a) values for each subject. The orange dot represents the mean and the blue line represents the HPD interval. (d) Illustrates the 95% HPD interval of the Non-Decision times values for each subject. The orange dot represents the mean and the blue line represents the HPD interval. (e) The violin plot represents the posterior distribution of the Drift Rate on the different days, for the different coherence levels, and for face stimuli (top) and car stimuli (bottom). Here, we use violin plots to illustrate the posterior distribution. Violin plots plot the probability density using a symmetric shape. To understand the violin plot, imagine a vertical line through the centre of each shape. The shape will be symmetric around this axis. The density on shown on either side of this axis is the density plot for the variable.

negative drift rates. As learning proceeds, we see that the absolute value of these drift rates continue. We also note how the absolute values of the drift rates increase when the stimuli have higher rather than lower coherence. These results demonstrate how efficiency of information accumulation increases with training and with the decreasing noise in the stimuli.

5.3 Diffusion model analyses of the effect of prior information on perceptual decision making.

In this analysis, we model response speed and choice data from our second project, i.e. the project described in Chapter 4, using a HDDM. Here, our particular aim is to determine how prior probability of the upcoming stimulus, as revealed by the pre-stimulus cue, affects perceptual decision making by examining which parameters of the diffusion model vary with experimental predictor variables, but particularly with the pre-stimulus cue.

In this model, as above, our outcome variables are the participant's choice, i.e. face or car, on each trial, and their reaction time to make the choice. These are modelled by first passage times in the DDM: Crossing the upper barrier signifies the choice of face, and the lower barrier signifies the choice of car. In the DDM, in general, the drift rate, inter-barrier distance, bias, and the non-decision time can all, in principle, vary by day of learning and by pre-stimulus cue each trial. Again, and as above, the model is a HDDM given that drift rate, inter-barrier distance, bias, and the non-decision time all vary randomly by subject.

Note that in this model, despite the fact that stimuli varied by coherence levels, we chose not to model the effects of coherence explicitly. From the previous analysis described above (for example, Section 5.2.1, page 108), it was clear that coherence had an effect only on the drift rate. In particular, when coherence was low, the absolute value of the drift rate was lower, and when the coherence was high, the absolute value of the drift rate was greater. This is perfectly in line with expectations, i.e., when there is less noise, there is a faster accumulation of evidence. And so it will undoubtedly be the case that coherence would have the same effect on the drift rate in the present model. On the other hand, coherence levels can not affect the bias, which is a diffusion model variable that we

are including here, because bias is only affected by variables that are present prior to the stimulus onset, such as pre-stimulus cue.

By not explicitly modelling coherence, we are not assuming that coherence plays no role in the model. As mentioned, we do assume that it will affect the drift rate. However, by not modelling it explicitly, we are effectively treating coherence level as a constant variable across all trials, and so modelling its average effect on drift rate. We are choosing not to model coherence explicitly in order to facilitate the analyses. Given that we have new model variables and new predictor variables in this current model, not explicitly modelling all variables allows us to keep the models tractable. Specifically, we have a new predictor variable, i.e., cue, and this will potentially affect the drift-rate and bias. In initial MCMC simulations, where we included all the principal variables, we obtained very poor convergence rates in the model. By not explicitly modelling coherence, however, convergence rates improved. Given that the primary focus of this current model is to identify how the cue and day of training affect either the bias or the drift rate, by not explicitly modelling coherence and treating it as a constant variable, we will not be affecting the general conclusions from this analysis. Thus, by not explicitly modelling coherence, we may proceed with the current more complex model analyses, and yet not affect the main conclusions about the role of the cue and day of training on the bias and the drift rate.

5.3.1 Model selection

As above (see Section 5.2.1), we begin our analysis by considering a range of possible HDDM models. In order to minimize the proliferation of models to consider, we restricted attention to two main predictor variables, i.e., day of learning and pre-stimulus cue. We also consider the interaction of these variables. In particular, we are interested in whether the effect of pre-stimulus cue on perceptual decision making changes with training. We then restricted our attention to how these two variables and their interaction affect the drift-rate and the bias. We consider all possible models where one or more predictor variable, possibly with an interaction, affect either drift rate or bias or both. This leads to 16 models. The WAIC for each model is shown in the following table:

WAIC	model description
3378.27	drift ~ day; bias ~ day
3234.82	drift ~ day; bias ~ cue
2939.01	drift ~ day; bias ~ cue + day
3449.46	drift ~ day; bias ~ cue + day + cue × day
4200.43	drift ~ cue; bias ~ day
3218.35	drift ~ cue; bias ~ cue
2928.38	drift ~ cue; bias ~ cue + day
3312.89	drift ~ cue; bias ~ cue + day + cue × day
3717.17	drift ~ day + cue; bias ~ day
2809.44	drift ~ day + cue; bias ~ cue
2680.16	drift ~ day + cue; bias ~ cue + day
3176.03	drift ~ day + cue; bias ~ cue + day + cue × day
3294.86	drift ~ day + cue + cue × day; bias ~ day
3553.38	drift ~ day + cue + cue × day; bias ~ cue
2734.23	drift ~ day + cue + cue × day; bias ~ cue + day
3362.32	drift ~ day + cue + cue × day; bias ~ cue + day + cue × day

As can be seen, the model with the lowest WAIC, and so the model we choose, shows that both bias and drift rate vary by day and cue. A graphical model illustration of this best fitting model is shown in Figure 5.5. Note that in all models, the drift rate, bias, boundary separation, and non-decision time are modelled as varying randomly by participant.

5.3.2 Analysis of best fitting model

As before, we begin by comparing the predictions of the model to the empirical data from the subject. In Figure 5.6a and Figure 5.6b, we provide scatterplots comparing model predictions to the subjects' reaction times, for both stimuli classes, pre-stimulus cue levels, and days of training, and for correct and incorrect reaction times. For the correct responses, the correlation coefficient for the fits to the face and car stimuli are as follows:

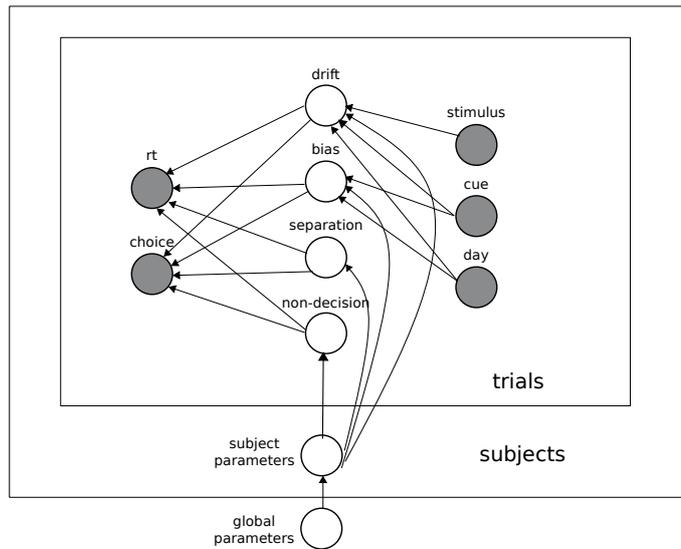


Figure 5.5: **A graphical model illustration of the best fitting model of the effect of prior information on perceptual decision making.** A graphical model illustration of the best fitting model for the data from the prior expectation experiment (see Chapter 4). This shows which predictors and other observed variables affect or influence which parameters in the system, and how these affect the accuracy and reaction time. The shaded nodes indicate observed values of the data.

Car	Face
0.87	0.83

For the *incorrect* responses, the correlation coefficient for the fits to the face and car stimuli are as follows:

Car	Face
0.78	0.80

In both cases, we can see that the models' reaction time predictions are relatively accurate. In all cases, we see that the models' predictions always resemble the empirical reaction for the different stimuli and on the different days.

As we did for the model for the first project, in Figure 5.6c and Figure 5.6d we present the 95% HPD intervals for the boundary separation parameter and non-decision times, respectively, for each subject. Here, as above, for the boundary separation variable, the HPD

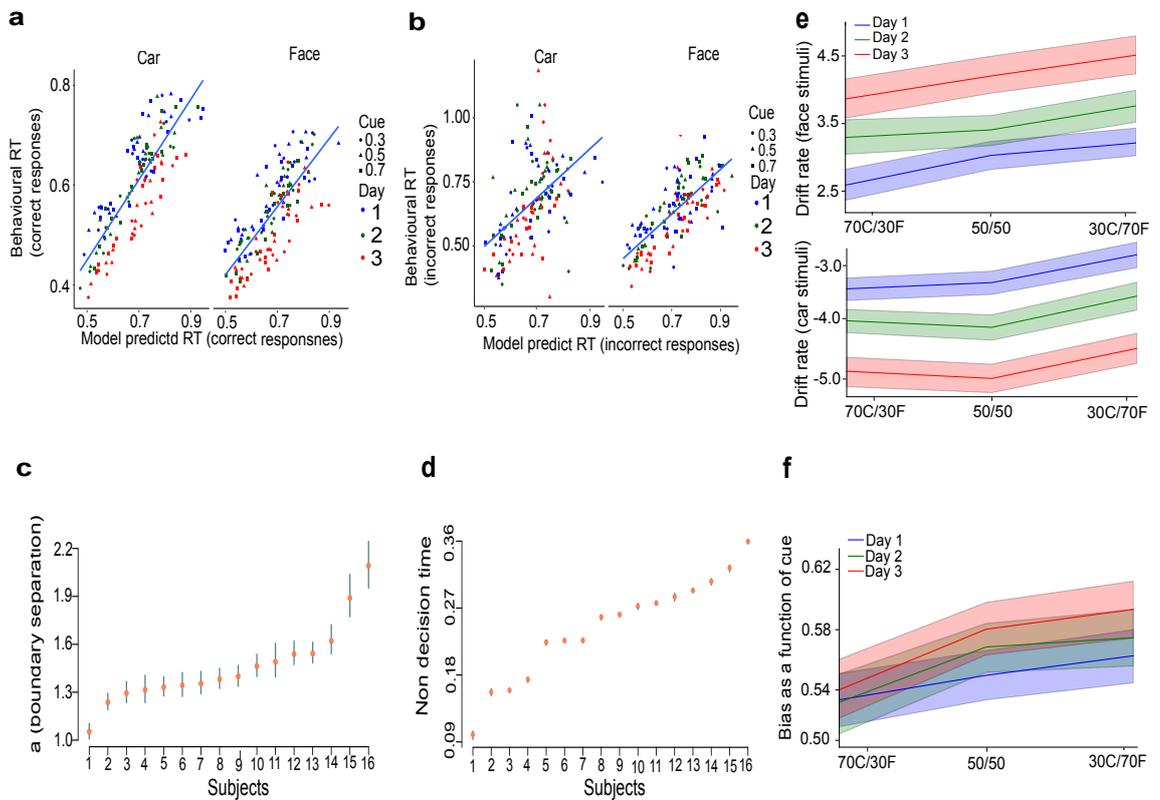


Figure 5.6: Diffusion Model of the effect of prior information on perceptual decision making.

(a) Illustration of model fit. The model’s predicted reaction times for stimuli on each day and for each pre-stimulus cue are compared to the average reaction times of all subjects for these variables. Here, we show reaction times according to the model and from the subject, where the responses were correct. (b) The model’s predicted reaction times on each day and for each cue are compared to the average reaction times of all subjects for these variables. Here, we show reaction times according to the model and from the subject where the responses were incorrect. (c) Shows the 95% HPD interval of the Boundary Separation (a) values for each subject. The orange dot represents the mean and the blue line represents the HPD. (d) Illustrates the 95% HPD interval of the Non-Decision times values for each subject. The orange dot represents the mean and the blue line represents the HPD. (e) The 95% HPD of the Drift Rate for both face (top) and car stimuli (bottom) as a function of pre-stimulus cue on each of the three days of learning. (f) The 95% HPD for the Bias variable as a function of cue on each of the days of learning (1: blue; 2: green; 3: red) across 16 subjects.

intervals show how subjects differ from one another in terms of their speed accuracy trade-off. For the non-decision time, the variability shows how subjects differ considerably in their sensory-coding and motor response times.

In Figure 5.6e, we show the 95% posterior HPD for the drift rate for face and car stimuli as a function of the different cue levels, and for each of the three days of learning. As can be seen, there is noticeable increase in the absolute value of the drift rate with each day of learning. In other words, for faces, the drift rate increases with training, and for cars, the drift rate decreases. As was the case for the analysis of the data from the first project, these results demonstrate how efficiency of information accumulation increases with training and with the decreasing noise in the stimuli. We also see here that as the cue indicates a higher probability of the upcoming stimulus, the absolute value of the drift rate increases. Thus, when as the cue progress from 30% probability of face to 50% probability to 70% probability, the drift rate for face stimuli likewise increases. For car stimuli, as the cue progresses from 30% probability of car to 50% probability to 70% probability, the drift rate for car stimuli becomes increasingly negative.

In Figure 5.6f, we show the bias variable varies with cue probability and on each day of learning. As can be seen, the bias increases towards a face response with the increasing probability of a face stimulus according to the cue. We also notice a gradual rise in the average face bias with each day of learning.

5.4 Discussion

The diffusion model analysis of the data from the first project, i.e. the project presented in Chapter 2 and in Diaz et al. (2017), reveals that the locus of the effect of learning on the increased speed and accuracy that we described in Chapter 2 is due to the increased absolute value of the drift rate with learning. Note that the model comparison analysis for the HDDM models revealed that the model fit when learning was included as a predictor of both boundary separation and non-decision times was in fact poorer compared to when these parameters were constant effects. This result alone is sufficient to establish that learning does not reliably affect either the boundary separation or non-decision time in the diffusion model. The only variable that learning does affect is the drift rate, and this

effect is clearly visible in the posterior plots shown in Figure 5.4e. This result largely corroborates results presented in Petrov, Van Horn, and Ratcliff (2011) who showed a stimulus-specific increase in the drift-rate parameter in a perceptual learning experiment over the course of multiple days. As such, our results provide a replication of some of the main findings of Petrov et al. (2011) and so provide additional evidence, along with the results presented in Chapter 2, confirming that the effect of learning on perceptual decision making is to increase the speed and efficiency of evidence accumulation.

In addition to establishing the main effect of learning is to increase the drift rate, this study corroborates results presented in Philiastides et al. (2006b); Ratcliff et al. (2009a). Philiastides et al. (2006b) showed a correlation between the strength of the late component and the mean drift rate in a diffusion model that was applied to six subjects at five different phase coherence levels. Ratcliff et al. (2009a) divided trials on a face/car discrimination task on the basis of the amplitude of the late temporal component. Doing so, they found that there was positive correlation between the late component amplitude and the drift rate in the diffusion model. Dividing the trials in terms of the amplitude of the early component resulted in no correlation between the amplitudes of the early component and the drift rate.

For the case of the data from the experiment concerning the role of prior expectation, the model fit analysis here was focused on how the day of learning and the cue probability affects the drift rate and the bias. The model fit statistics clearly show that both the drift rate and the bias are independently affected by the day of learning and the cue probability. In neither case do we see any evidence in favour of an interaction of these two predictor variables. While the model fit statistics are sufficient to establish a statistically reliable effect of cue probability on both the drift rate and bias, the effect on drift rate is both clearer and more straightforward. In other words, as the probability of the upcoming stimulus increases, so too does the absolute value of the drift rate. This shows that prior expectation as indicated by the pre-stimulus cue leads to a faster and more efficient accumulation of evidence when making a decision. These results are consistent with Hanks et al. (2011); Cravo et al. (2013); Dunovan et al. (2014), reviewed in Chapter 4, who argued, using results from NHPS and human subjects, that the prior expectation increases the accumulation of evidence to reach a decision threshold.

In comparison to the effect of cue probability on drift rate, the effect of cue on bias is less clear. While we observe that with increasing probability of a face response, there is an increase in the bias towards faces, we do not observe an increasing bias towards cars when the probability of car stimulus is high. In other words, there is always a bias in favour of face stimuli and this increases as cue predicts face stimuli. This average bias towards the face responses does not have an obvious explanation. Moreover, the rise, albeit limited, in this average bias towards faces as training increases is also an anomalous result.

Comparing our results for these diffusion model simulations with those Dunovan et al. (2014), the results are similar with respect to the drift rate. In both Dunovan et al. (2014) and here, a clear result of drift rate increase with training is observed. However, Dunovan et al. (2014) observed a much clearer effect of learning on bias. In our analysis, the effect of learning on bias was less strong. It is worth noting as well that in the finding presented in Philiastides et al. (2006b); Ratcliff et al. (2009a), the correspondence between the EEG late temporal component is also observed only with respect to the drift rate.

5.4.1 Investigating the possible effects of outlier

One possible explanation of the relatively unexpected findings with respect to the bias may have been due to an extreme influence caused by outliers. In particular, the distribution of reaction times necessarily assigns a probability density of zero to any reaction time that occurs in the non-decision time interval. When fitting the model, therefore, the non-decision time must be *at most* equal to the fastest reaction time. Occasionally, extraordinarily fast reaction times may occur by accident, for example, by a premature key-press almost at the moment the stimulus presentation. Whenever, these reaction times occur, they will necessarily lead to the non-decision time being reduced to an unrealistic value. As a consequence, estimated value of all other parameters may also be affected.

In order to consider this possibility further, we plotted all reaction times by all participants in our two data-sets. This is shown in Figure 5.7. As can be seen, in both experiments, the distribution of the data is typical of the distribution of reaction times. In particular, there is a considerable positive skew, i.e., a long tail of reaction times in the positive/higher direction. Although some of these data will necessarily be far from the

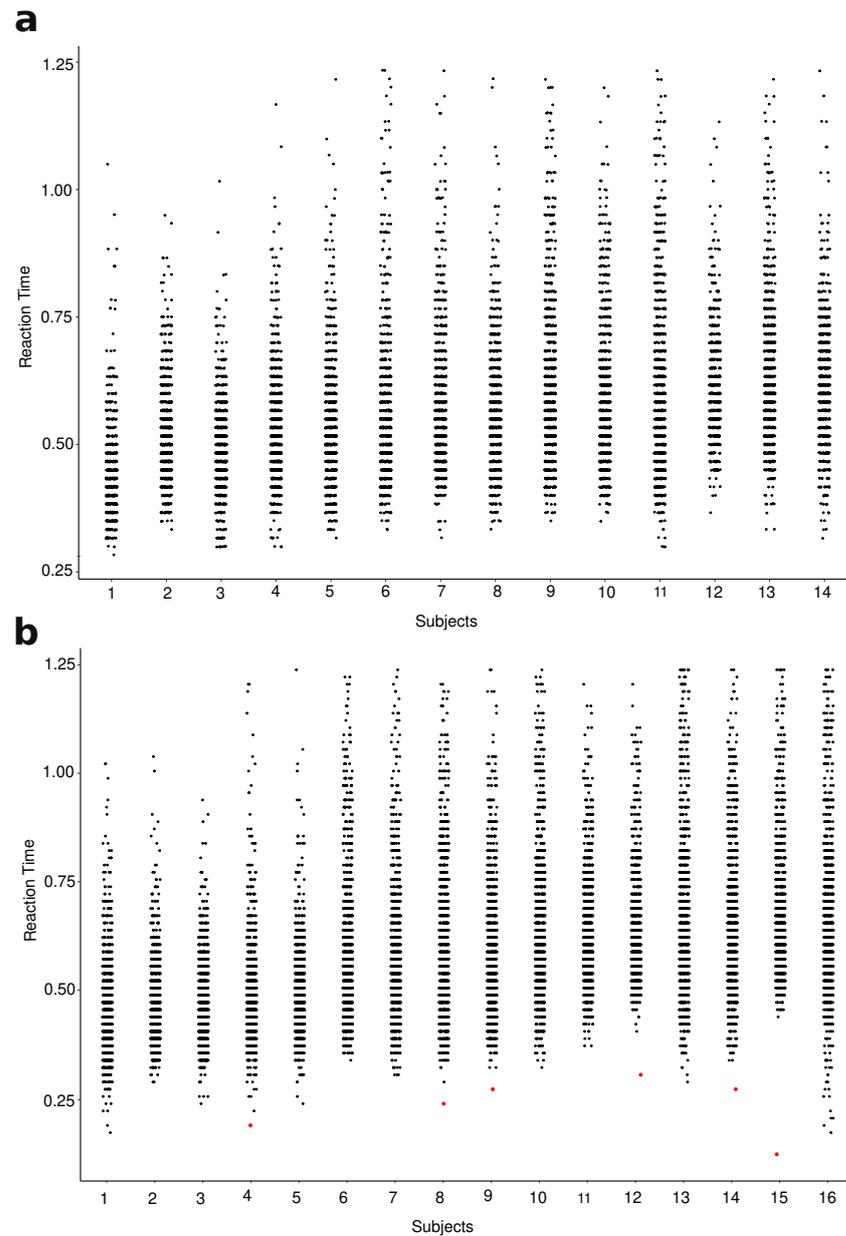


Figure 5.7: **Plots of all reaction times to identify outliers.** Here, we plot every reaction time by each subject in each of our two experiments (**a** shows data from Experiment 1, **b** shows data from Experiment 2). In all cases, the distribution of points is typical of reaction time data. In particular, we see a considerable positive skew, i.e. there is a long tail of reaction times in the positive or higher direction. None of the points in the positive tail can be regarded as outlier per se. On the opposite tail, in the data from Experiment 2, there are a small number of points, shown in red, that can be regarded as legitimate cases of outliers.

mean, they do not constitute outliers per se. An outlier is defined generally as a point that does not conform to a particular distribution.

For example, in a normal distribution, given that only around one half of 1 percent of data will be beyond around 2.5 standard deviations from the mean, it is reasonable to regard points much greater than this as outliers. In other words, any points much greater than this threshold can be regarded as not having arisen from the normal distribution that describes the remainder of the data. Using the same reasoning, however, we can not easily regard any points in the upper tail of the distributions depicted in Figure 5.7 as outlier because, by the nature of the highly skewed probability distribution of reaction times, there will necessarily be a large number of points sparsely spread out in the positive direction. Moreover, for the immediate practical problem we are investigating, as mentioned above, it is fast reaction times are the potential source of anomalous results, and so in this respect, large reaction times are not an immediate problem.

Turning our attention to the lower tail, there are no obvious cases of outliers in the data from Experiment 1, which are depicted in Figure 5.7a. In particular, there are no points that are more extreme than approximately 15ms faster than others. In other words, across all subjects in this experiment, the very fastest reaction times are within 15ms lower than the next fastest times. Given that the resolution of the reaction time recording is approximately 15ms³, there is virtually no gap between the very fastest reaction times and the next fastest ones, and the fastest reaction times are not anomalous and should not be regarded as outliers.

In the data from Experiment II, i.e., the data depicted in Figure 5.7b, there are clearly some outliers. Most noticeably there is the one extremely fast reaction time, hundreds of milliseconds faster than the next fastest, in the case of Subject 15. Defining as an outlier any reaction times faster by more than 30ms from the next fastest reaction times, we can identify seven outliers. These are depicted in red in Figure 5.7b.

There are two important points to make out these outliers. First, there are very few of them. Second, with the exception of perhaps a single point, they are not extreme. As such, it is unlikely that their removal from the data will be of any practical consequence.

³PsychoPy uses the graphics card rate to record reaction times. We use a 60Hz graphics card, so the resolution is close to 16ms.

In order to test this, having removed these outliers, we then re-ran all 16 models, and re-calculated the WAIC for each. As expected, there were no major differences in the pattern of results, and in particular, the same models were identified as being the best fitting ones. From this, we conclude that outliers are unlikely to have caused any adverse effect on the main results reported in this chapter. In particular, we were concerned with any extremely fast reaction times adversely affecting the parameter estimates. Because there may have been only a single extra fast reaction time data point, with the other points that we have identified as outliers being neither extremely fast, nor very much faster than other reaction times by those participants, it is implausible that the presence of these few data points in data-sets of tens of thousands are likely to have had any major effect on the results.

Chapter 6

General conclusions

In psychology and neuroscience, decision making can be generally described as the commitment to a classification of data, which is then usually followed by a voluntary response or action. A special case of decision making is perceptual decision making, which is the classification of sensory, or sensory-motor, information usually followed by an overt behavioural response. In any perceptual decision making, we experience auditory, visual and other sensory information, classify that information as belonging of some category, and then respond accordingly.

All decision making can be described from a statistical point of view: There is observed data, a set of possible classifications of this data, and a reward or loss function that gives the payoff or loss for any given choice. Bayes's rule allows us to calculate the probability that a given classification of the observed data is correct, and we then combine that with the loss function to determine the expected reward or punishment for any given decision that we make. This time course of this process of weighing evidence and making a choice when the evidence is sufficient to indicate a high probability of a successful classification can be modelled using sequential sampling models, such as the DDM (e.g. Ratcliff, 1978; Ratcliff & Smith, 2004; Ratcliff et al., 2004; Ratcliff & McKoon, 2008; Ratcliff & Van Dongen, 2011). Sequential sampling models have proven effective in modelling both the behavioural and neuroscientific aspects of decision making. In human neuroscience using EEG recording, studies such as O'Connell et al. (2012); Kelly and O'Connell (2013) and Philiastides and Sajda (2006); Philiastides et al. (2006a, 2006a);

Ratcliff et al. (2009b) have proven effective in measuring the nature and time-course decision making in the human brain.

In this thesis, our aim has been to address some open questions with respect to human perceptual decision using the theoretical framework of sequential sampling models and the experimental paradigm of measuring temporal components in single-trial EEG discriminant analysis. As we have laid out in Chapter 1, the two open questions that we will address in this thesis are how does learning or training affect perceptual decision making, and how do prior experiences interact with learning to affect perceptual decision making.

6.1 The neural locus of perceptual learning

In Chapter 2, we provided evidence that perceptual learning arises from changes in higher level brain areas that are related to decision-making, rather than from perceptually earlier areas that are related to the encoding of sensory stimuli. Our results are premised on previous single trial EEG analyses (Philiastides & Sajda, 2006; Philiastides et al., 2006a; Ratcliff et al., 2009b) that have established that there are two major EEG temporal components that correlate with perceptual discrimination behaviours: One component, the *early* component, has been shown to be indicative of sensory processing, while the second *late* temporal component has been shown to be the neural signature of the decision-related part of the perceptual classification task. In Chapter 2 we have shown that the later EEG temporal component but not the early one increases in strength and occurs earlier in time as a function of perceptual learning. In particular, we show that for each extra day of perceptual training, the area under the ROC curve (a value bounded between 0.5 and 1.0), which quantifies the discriminating performance of the EEG components, significantly increases for the late component but does not significantly change for the early component. Likewise, the onset time of the late component occurs significantly earlier with every day of training, while there is no significant change in onset time for the early component.

The primary significance of these findings lies in how they relate to the ongoing debate about the neural locus and information processing stage at which perceptual learning occurs. In particular, given that perceptual learning is inextricably related to perceptual decision making — improvements in detecting, discriminating, identifying or respond-

ing to sensory stimuli all necessarily involve improvements in the speed and accuracy of perceptual decision making — an open question is the extent to which perceptual learning is due to improvements in general decision making processes or due to improvement in sensory abilities that are (informationally and temporally) prior to perceptual decision making.

The traditional perspective on the nature and locus of perceptual learning is that it arises early in the perceptual system. For example, well known seminal results from psychophysics have demonstrated that perceptual learning is often highly specific to the location and other properties of the stimuli (Fiorentini & Berardi, 1980; Poggio et al., 1992; Fahle & Edelman, 1993; Karni & Sagi, 1991; Crist et al., 1997; Ball & Sekuler, 1987; Sagi & Tanne, 1994), implying that its plasticity is highly specific to the trained retinal location (Fahle, 2004, 2005). In contrast to this traditional perspective, alternative psychophysical studies have proposed that perceptual learning arises not from plasticity in primary sensory areas, but rather from changes in how sensory signals are read out or interpreted by decision-making mechanisms Doshier and Lu (1999); Petrov et al. (2005); Lu et al. (2010). Indeed, results from NHPs (Law & Gold, 2008a, 2009) have provided compelling evidence in support of this general hypothesis, demonstrating that perceptual learning involves areas beyond the perceptual system, and in particular, decision making centres of brain such as LIP and ACC.

The results that we have presented in Chapter 2 provide corroborating evidence for the results obtained by (Law & Gold, 2008a) and show that their findings with primates also extend to the human brain. In particular, we provide evidence in humans that perceptual learning does not change how sensory information is *represented* in the brain, but rather how sensory representations are *interpreted*, particularly by higher areas in the brain involved in decision making. For example, the finding that the late decision-related component increases as a function of learning and systematically moves backward in time is consistent with a faster and more efficient accumulation of evidence and this is indicative of the strengthening of the connections between early sensory encoding regions and the final areas involving choice commitment.

6.2 Reinforcement learning models of perceptual learning

Previous accounts of perceptual learning by Law and Gold (2009); Kahnt et al. (2011) have showed that learning-induced behavioural improvements in perceptual decision making could be reliably explained in terms of a reinforcement learning (Sutton & Barto, 1998) mechanism. The behavioural and EEG results we present in Chapter 2 are also compatible with this account. In Chapter 3, on the basis of a reinforcement learning model identical to Kahnt et al. (2011), we described a highly significant correlation with the behavioural error rates in the perceptual task. More importantly, we have established that the reinforcement learning model's decision variable corresponds more closely to the y variable measured at the late component than at the early component. These results imply that reward based feedback from decisions may play a vital role in the trial by trial learning in perceptual tasks.

In addition, as described in Section 3.4, in Diaz et al. (2017), we provided an RL model that is an extension of our initial model. The main innovation of this extended model is that it includes both signal and noise weights. While the signal weights were designated to enhance stimulus representations, the noise weights accounted for the interference exerted by the antagonistic stimulus against the acquisition of the correct sensory associations. Thus, in this model, perceptual learning is expected to occur through gradually increasing signal weights as well as gradually decreasing noise weights. Compared with Law and Gold (2009); Kahnt et al. (2011), this better captures instances whereby improved task performance depends both on greater ability to recognize a given stimulus and on greater ability to rule out the antagonistic stimulus. For example, on a face trial, subjects might correctly choose the face response either because they are able to identify face-like features or else because they are able to recognize that there are no car-like features, or else through some combination of both these two pattern recognition processes.

6.3 The role of prior expectation in perceptual learning

In Chapter 4, our behavioural results clearly demonstrate that as the probability of the upcoming stimulus according to the cue increases, so too does the speed and the accuracy of the perceptual decision. On the basis of a discriminant analysis identical to that described in Chapter 2, we also showed that how the pre-stimulus cue affects the late temporal component parallels how the cue affects the probability of making a face response. Notably, we did not find a similar correspondence between the probability of a face response and the effect of the cue on the early component. The results in this chapter are consistent with how prior expectation affects evidence accumulation, as revealed by, for example, Hanks et al. (2011); Cravo et al. (2013); Dunovan et al. (2014). While these results are less consistent with the baseline activation hypothesis as forwarded by, for example, Basso and Wurtz (1998); de Lange et al. (2013); Albright (2012), we did observe a post-cue pre-stimulus EEG component that is arguably related to attentional mechanisms. This latter result may imply that pre-stimulus cues also lead to a change in baseline activation of sensory processes.

6.4 Drift diffusion models of perceptual learning

The diffusion model analysis of the data from Chapter 2 reveals that the locus of the effect of learning on the increased reaction time and increased accuracy that we described in that chapter is largely due to the increased absolute value of the drift rate with learning. The model comparison analysis established that learning does not reliably effect either the boundary separation or non-decision time in the diffusion model. The only variable that learning does effect is the drift rate. These results are in line with results by Petrov et al. (2011) who showed an increase in the drift-rate parameter in a perceptual learning experiment. They also corroborate results presented in Philiastides et al. (2006b) that showed a correlation between the strength of the late component and the mean drift rate in a diffusion model, and by Ratcliff et al. (2009a) who found that there was positive correlation between the late component amplitude and the drift rate in the diffusion model.

For the case of the data from the experiment described in Chapter 4, the HDDM simulations show that both the drift rate and the bias are independently affected by the day of learning and the cue probability but we do not see any evidence in favour of an interaction of these two predictor variables. We show that as the probability of the upcoming stimulus increases, so too does the (absolute) value of the drift rate, and this implies that prior expectation leads to a faster and more efficient accumulation of evidence when making a decision.

6.5 Limitations, extensions, & future directions

Collecting more behavioural data: The behavioural data that we collected in our main experiments was collected while the participants were being EEG recorded. Because collecting EEG data is resource and time expensive, it was not possible to collect very large amounts of data. However, if we were to restrict ourselves solely to the collection of the behavioural data alone, it would be possible to collect larger amounts of behavioural data. While behavioural data alone is obviously not sufficient to address our main questions, larger amounts of behavioural data could be particularly valuable for some questions. For example, all of the analyses using the DDM and HDDM described in Chapter 5 were based solely on behavioural, and not EEG, data. Had we performed these experiments with larger numbers of participants, we would have been able to perform improved analyses. In particular, larger amounts of behavioural data would have been able to provide more certainty concerning the parameters of the HDDM. It would have also provide more statistical power to allow us to better discriminate between the competing models that we used. As it currently stands, on the basis of the model comparisons, some of the predictor variables did not appear to affect some of the HDDM parameters. However, this may have been a consequence of lack of statistical power due to the relatively low numbers of participants. The general problem of lack of statistical power in neuroscience, particularly cognitive neuroscience, studies has been recently described in Button et al. (2013).

Alternative behavioural tasks and EEG analyses: An alternative EEG experiment that could be used to investigate the temporal processes of perceptual decision making is based on the experimental paradigm described in Kelly and O'Connell (2013); O'Connell et al.

(2012). Using a continuously flickering, but gradually changing visual stimulus, it has been claimed that this method has the potential to identify a supra-modal *accumulation-to-bound* decision making variable that manifests itself a CPP ERP. Although we have performed pilot experiments using this method (not described in this thesis), we have not fully explored the potential of this method to provide a complementary experimental paradigm for the study of the temporal dynamics of perceptual decision making.

Extensions to developmental studies: We have investigated how learning affects perceptual decision making. An important related question is the developmental trajectories of the perceptual decision making system. In other words, how does the temporal dynamics of perceptual decision making change from childhood to adulthood. This topic is related to studies of development changes in decision making studied using DDM models (e.g., Ratcliff, Love, Thompson, & Opfer, 2012). In principle, the main experimental paradigm that we have used, using a similar visual discrimination task, and modelled and analysed using a combination of single trial EEG analysis and HDDM computational models could be equally well employed using children and ages between childhood and adulthood.

6.6 Conclusion & main contribution to the field of perceptual learning and decision making

The present thesis presented the empirical findings and computational modelling from two major studies that investigated the neural signatures of perceptual decision making. The aim was to address two open question concerning human perceptual decision making: how does learning affect perceptual decision making and how do prior expectations interact with learning to affect perceptual decision making.

Our first study provided insights into the neurobiology of perception learning by showing that the locus of learning can be in higher level brain areas that are related to decision-making rather than necessarily arising in areas that are related to the encoding of sensory stimuli. Specifically, this study showed that perceptual learning involves larger changes in decision related stages of perceptual processing compared to the sensory processing stages. Moreover, it showed that these decision-related stages increase in importance as

a function of learning and systematically occur earlier in time, which is consistent with a faster and more efficient accumulation of relevant decision evidence. This particular result is also corroborated by drift diffusion modelling of the behavioural data, which shows a clear effect of learning on the efficiency of the accumulation of evidence for decision making. It also showed that behavioural improvements as a function of perceptual learning can be explained with a simple reinforcement learning model. This result has implications of our understanding of the relationships between general reward-based learning and decision-making mechanism in perceptual learning. Overall, this study provides critical insights into the neurobiology of perceptual learning and offers strong support to the notion that neuronal plasticity can occur at multiple time-scales and locations, depending on task demands and context. As such our findings can help revise existing theories of perceptual learning focusing only on early sensory processing and provide the foundation upon which future studies continue to interrogate the neural systems underlying perceptual decision making (see, Diaz et al., 2017).

From our second study, we have provided insights into the neural signatures of the effect of prior expectations on perceptual decision making in a learning task. This study both corroborates and extends our first study. We showed that there is a relationship between the predictive power of pre-stimulus cues and the strength of the decision related stages, but not the sensory processing stages, of perceptual learning. Also in this study, the drift-diffusion modelling shows that prior expectation affects both drift rate and the bias, which means that the pre-stimulus information can lead to a faster and more efficient accumulation of evidence during perceptual processing.

In summary, this thesis has provided compelling evidence that perceptual learning alters post-sensory processing in human decision-making and likewise that how prior expectations affect decision making is through its effect on post-sensory processing. Taking together, our results represent a step toward a more theoretical and experimental framework of the study of the neural signatures of human decision making in the context of perceptual learning.

Appendices

Appendix A

Pilot experiments from Chapter 2

A.1 Overview

Before embarking on the main experiment, we carried out a set of 5 separate pilot experiments. Of the 5 pilot experiments, 2 form what we call below *Pilot experiment set 1*, and two form *Pilot experiment set 2*. We call the fifth experiment, the *Main pilot experiment*. The main purpose of these pilot experiments is to verify if the accuracy and the speed of response increases over days, and whether these accuracy and speed effects increase with learning, and finally to see if there is any significant effect of the role of feedback on learning. In total, 34 subjects were used in these pilot experiments and each subject was tested on each of three consecutive days, giving a total of 102 experiment sessions. In all pilots, the stimuli and the general behavioural paradigm were identical.

Performing pilot studies is generally a necessary step to obtain a good experimental design. They provide us with vital insight into how the eventual main experiment is likely to proceed. Of particular importance in these pilot studies, however, was to investigate experimental settings, especially the ideal level of noise to use with the stimuli. Ultimately, we were looking for stimuli and noise levels that would lead to the clearest indication of perceptual learning in our subjects, as seen in both the subjects' speed and accuracies. Thus, in the different pilot experiments, the main difference between them lies in the number of distinct noise levels and the actual values of these noise levels.

A.2 Materials & methods

A.2.1 Stimuli

The stimuli used for this main experiments were drawn from the same set of stimuli described in the pilot experiments description in Chapter 2.

A.2.2 General behavioural paradigm

Trials in the pilot experiments were presented in 4 blocks of either 80 or 120 trials each. On each trial, a blank screen was displayed for a random duration that ranged uniformly between either 1100 and 1600 milliseconds or 1000 to 1500 milliseconds. The image stimulus was then presented for a fixed duration of 100 milliseconds. A blank screen was then redisplayed. Subjects were given up to 1250 milliseconds to make the response, which they did using the arrow keys (\Leftarrow for *face*, \Rightarrow for *car*). After a response, or after the timeout interval had elapsed, the experiment proceeded to the next trial. In some cases, as explained below, feedback after each trial was provided. Each experiment session was completed within approximately 20 minutes. Each subject performed this task on three consecutive days. They did so at approximately the same time of day on each day so that there was approximately 24 hours between each session.

A.2.3 Pilot experiment set 1

In the first set of pilot experiments, 3 levels of stimulus coherence were used. For one group of 5 subjects, the coherences levels were levels 25%, 30% and 35% for face images, and levels 30%, 35% and 40% for car images. For a second group of 6 subjects, the coherences levels were levels 27.5%, 32.5% and 37.5% for faces, and again levels 30%, 35% and 40% for cars. The two image categories, each with 20 different instances at 3 different levels of coherence gave $2 \times 20 \times 3 = 120$ unique images used as stimuli in these experiment. Each image was presented 4 times in total, giving 480 trials per session. These trials were presented in 4 blocks of 120 trials each, with a rest period of 60 seconds between blocks. The task design of this pilot experiment set, and pilot experiment set 2,

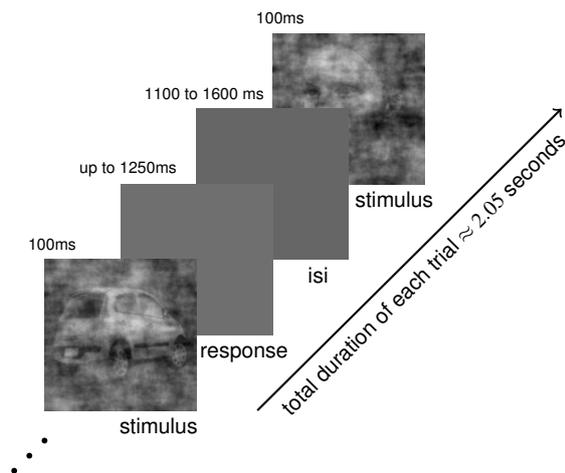


Figure A.1: The task design pilot experiment sets 1 & 2. The subjects' task was to indicate whether the briefly shown noisy image was either a *face* or *car*. Responses had to be made within the allowed interval of 1250ms.

are shown in Figure A.1. Descriptions of the stimuli of this pilot experiment set, and pilot experiment set 2, are shown in Figure A.2.

A.2.4 Pilot experiment set 2

In the second set of pilot experiments, 2 levels of coherence were used. For one group of 6 subjects, the coherences levels were levels 27.5% and 32.5% for faces, and levels 30% and 35% cars. For a second group of 6 subjects, coherence levels 30% and 35% were used for both face and car images. With two levels of coherence, this gave $2 \times 20 \times 3 = 80$ unique images, with each being presented 4 times for a total of 320 trials per experiment session. These were presented in 4 blocks of 80 trials, with a 60 second rest period between blocks.

A.2.5 Main pilot experiment

The main experiment used 2 levels of coherence. Coherence values of 32.5% and 37.5% were used for both faces and car images. The number of trials, blocks and trials per block was identical to that of pilot experiment set 2. In this experiment, all details of the general procedure remained the same as in the other pilot studies with the exception that now

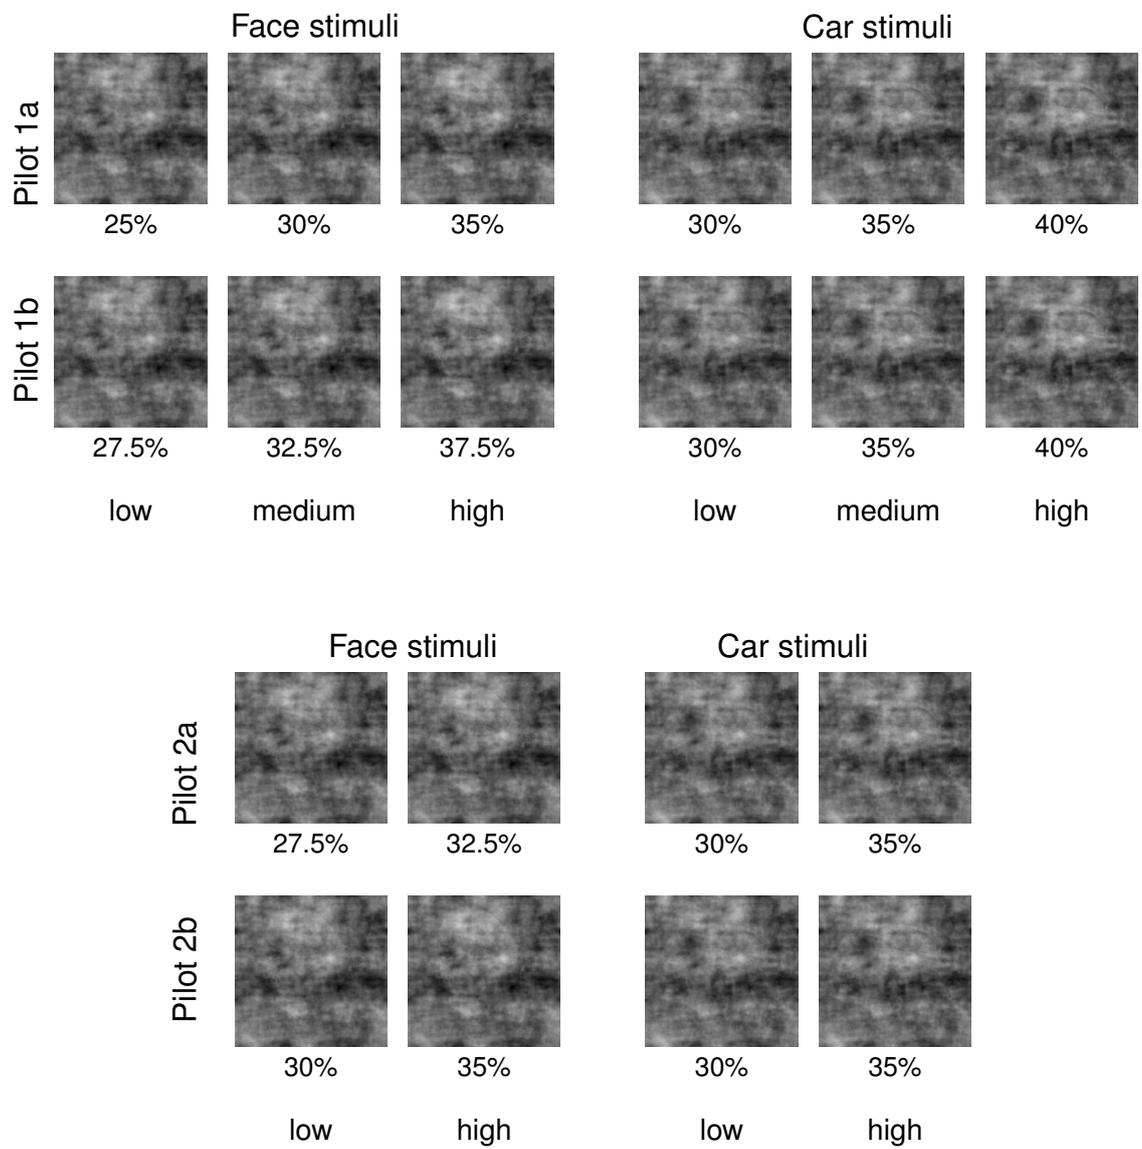


Figure A.2: Examples of stimuli used in Pilot Experiment 1 (above) and Pilot Experiment 2 (below).

feedback on responses was provided and the durations of the intervals before and after the stimulus presentation were modified to minor extents. Now, at the start of each trial, a blank screen was displayed for a random duration that ranged uniformly on the interval 1000 to 1500. The stimulus was displayed, as previously, for 100 milliseconds and a interval of 1250 was allowed for a response. If the subject responded correctly during this interval, a ✓ image was displayed for 500 milliseconds. If the subject was incorrect, or if they did not respond within the allowed interval, a ✗ was displayed for 500 milliseconds. The task design and the description of the stimuli for the main pilot experiment are shown in Figure A.3.

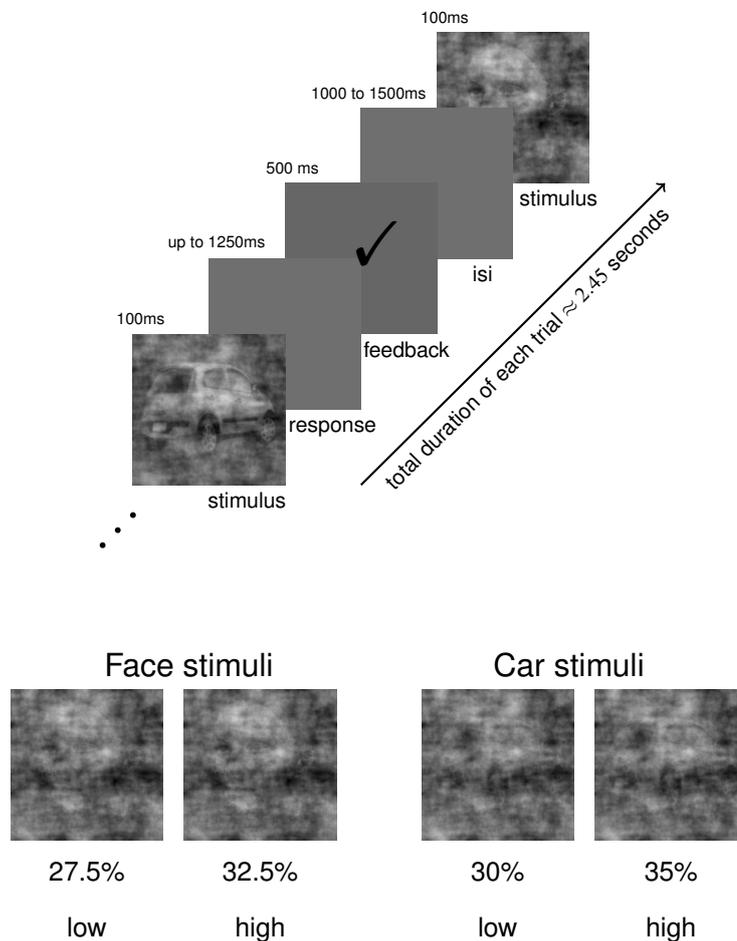


Figure A.3: The task design and stimuli for the main pilot experiment. The principal difference between the task here compared to that of the other pilot experiments is that feedback after each trial was provided.

A.3 Results

Reaction time and accuracy for pilot experiment sets 1 and 2 and the main pilot experiment are shown in Figure A.4. The data used in these plots was obtained as follows. Trials in which the subjects did not respond in the response interval — 1.5% of trials in pilot experiments set 1, 2.8% of trials in pilot experiment set 2, 1.5% of trials in the main experiment — were excluded from the analysis. Although the actual coherence values differed across the experiments in each set, we treated the levels in these experiment as effectively *high*, *medium* and *low* for the three level experiments and *high* and *low* for the two level experiments. Using the *hit* trials, we calculated the average accuracy on each day and for each of the levels of coherence, averaging over all trials and all subjects. For trials on which the response was accurate, we also calculated average reaction time latency for each day and each coherence level, again averaging over all trials and all subjects.

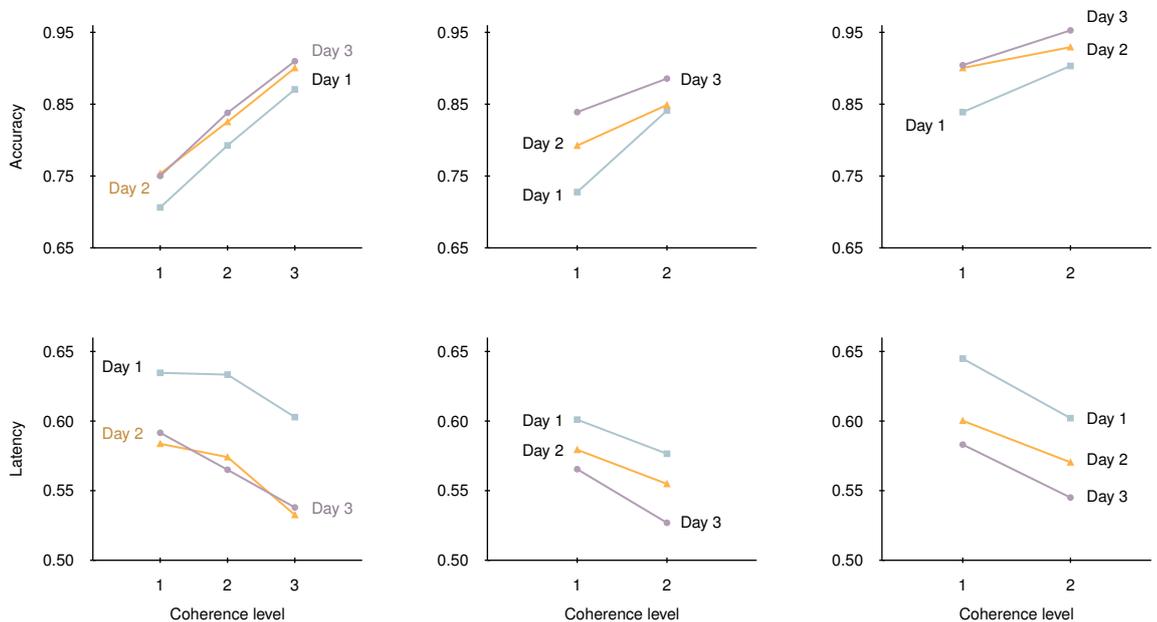


Figure A.4: Accuracy (top row) and reaction time latency (bottom row) for pilot experiment set 1 (left column), pilot experiment set 2 (middle column) and the main pilot experiment (right column).

As can be seen in Figure A.4, there is a clear apparent tendency for accuracy to increase and for reaction time latency to decrease over the three testing days. This is evident in both sets of pilot experiments and in the main experiment. This general trend manifests itself with minor variations across each experiment version.

As mentioned, we modelled the accuracy data using a mixed effects logistic regression analysis, and the reaction time latencies for accurate responses using a mixed effects linear regression. The level of coherence (i.e. *low*, *medium*, *high* for pilot set 1; *low* and *high* for pilot set 2 and the main pilot experiment) and day of session (i.e. *day 1*, *day 2* and *day 3*), and the interaction of coherence level and day of session, were treated as fixed effects in this analysis. The identity of the prototype image (i.e. for each stimulus, which of the 20 original faces or 20 original cars was used), and the identity of the subject were treated as random effects. Treating image prototype identity and subject identity as random effects allows us to model possible inter-item and inter-subject variability in the effects, and does so while avoiding the known inadequacies of *repeated-measures* analyses (see, e.g., Baayen et al., 2008, for further discussion). For each analysis, we used a log-likelihood ratio test — which can be formulated using a χ^2 test — to test for main effects and interaction effects.

For the accuracy analysis, in both sets of pilot studies and in the main experiment, there were highly significant main effects for both the level of coherence and the day of the session. There were only marginally significant effects of an interaction between the two main variables. The values of the χ^2 statistics — the degrees of freedom are shown in the parenthetical subscript — and their p-values are shown in the following table, where the rows represent the study and the columns represent the main or interaction effects.

	coherence	session	interaction
Pilot set 1	$\chi^2_{(2)} = 57.43, p < 0.01$	$\chi^2_{(2)} = 52.92, p < 0.01$	$\chi^2_{(4)} = 2.88, p = 0.06$
Pilot set 2	$\chi^2_{(2)} = 13.29, p < 0.01$	$\chi^2_{(2)} = 76.61, p < 0.01$	$\chi^2_{(2)} = 5.97, p = 0.05$
Main pilot experiment	$\chi^2_{(2)} = 9.56, p < 0.01$	$\chi^2_{(2)} = 160.50, p < 0.01$	$\chi^2_{(2)} = 4.97, p = 0.08$

For the latency analysis, a very similar pattern was observed to that found in the accuracy analysis. In both sets of pilots and in the main experiment, highly significant main effects of both level of coherence and day of session were found. However, no

significant effect of an interaction between these two variables was found. The values of the χ^2 test statistic and their p-values are shown in the following table.

	coherence	session	interaction
Pilot set 1	$\chi^2_{(2)} = 35.86, p < 0.01$	$\chi^2_{(2)} = 441.85, p < 0.01$	$\chi^2_{(4)} = 6.92, p = 0.14$
Pilot set 2	$\chi^2_{(2)} = 6.95, p < 0.01$	$\chi^2_{(2)} = 53.90, p < 0.01$	$\chi^2_{(2)} = 1.52, p = 0.47$
Main pilot experiment	$\chi^2_{(2)} = 8.37, p < 0.01$	$\chi^2_{(2)} = 365.49, p < 0.01$	$\chi^2_{(4)} = 3.64, p = 0.16$

A.4 Discussion of pilot studies

On the basis of the pilot experiment results, we can conclude that the behavioural paradigm appears to be robust and working according to expectation. In all variants of the experiment, we have observed significant increases in both accuracy and response speed of decision making over the course of the learning period. Given that learning has occurred and that this decision making task has been shown to yield strong neural signatures (e.g., Philiastides et al., 2006b; Philiastides & Sajda, 2006; Ratcliff et al., 2009a), we can be confident that EEG-detectable learning based changes in decision making system will occur in the main experiment described next.

Appendix B

Pilot experiments from Chapter 4

B.1 Overview

Before embarking on the intensive main experiment, we carried out 3 pilot experiments to verify that the main behavioural phenomenon that we expect to observe will occur with our choice of experimental parameters. With 3 pilots, we will verify if the accuracy and the speed of response increases with the increasing match between the pre-stimulus cue and the observed stimulus category, and whether these accuracy and speed effects increase with learning.

In total, 18 subjects were used in these pilot experiments and each subject was tested on each of one or two consecutive days, giving a total of 24 experiment sessions. In all pilots, the stimuli and the general behavioural paradigm were identical.

As was the case for the pilot experiments described in Appendix A, the general aim of these pilot studies is to determine optimal experimental settings to observe the phenomenon we aim to study in the main experiment. Specifically, we are interested in determining the optimal levels of noise to use with our stimuli, the amount of training that is required to see the effect of prior expectations, and which type of pre-stimulus cues are optimal (for example, shapes or labels).

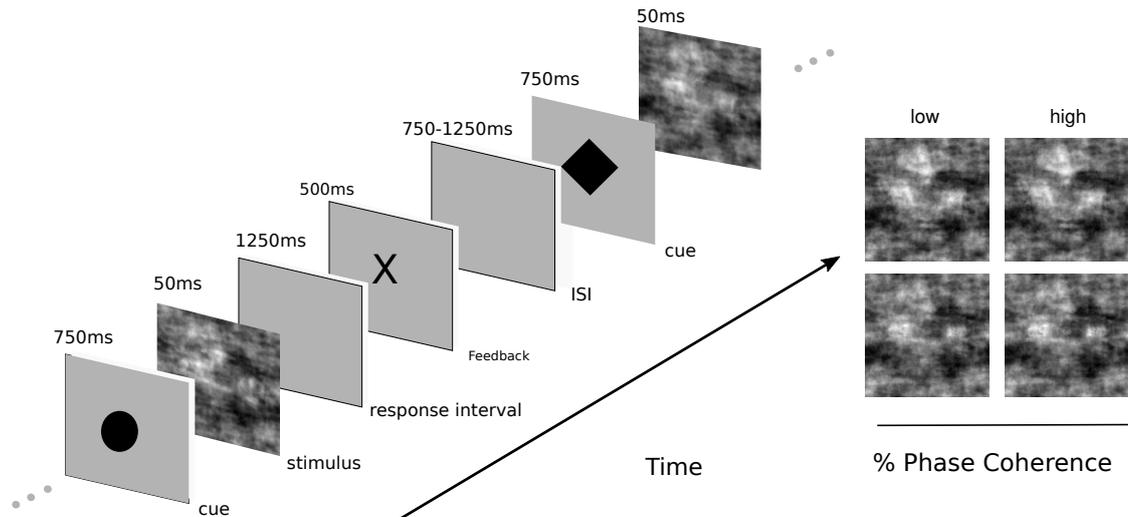


Figure B.1: The general characteristics of the behavioural task for the pilot experiments in Chapter 4.

B.2 General experimental design

For the three pilot experiments that we describe below, we always use the behavioural task shown in Figure B.1. On each trial of the experiment, a subject is shown a pre-stimulus cue, and then they must decide if a noisy image contains a face or car. The images were identical to the set of 18 noisy face and car stimuli that were used in the experiment described in Diaz et al. (2017) and in Chapter 2 of this thesis.

In all three pilots, there were three types of cues in the experiment: triangle, diamond, and hexagon. The triangle indicates the probability of a face stimulus being shown on the upcoming trial is exactly 0.30, which also means the probability of the upcoming stimulus being a car is 0.70. The hexagon cue has the opposite meaning, i.e., it indicates that the probability of a face stimulus being shown on the upcoming trial is exactly 0.70, and the probability of the upcoming stimulus being a car is 0.30. The diamond cue indicates the probability of the upcoming stimulus being a face or car is equally likely, i.e. probability is 0.50 for both stimuli.

In the different pilot experiments, we use different values for the low and high coherence, see below for details. As in the experiment described in Chapter 2, we also provide

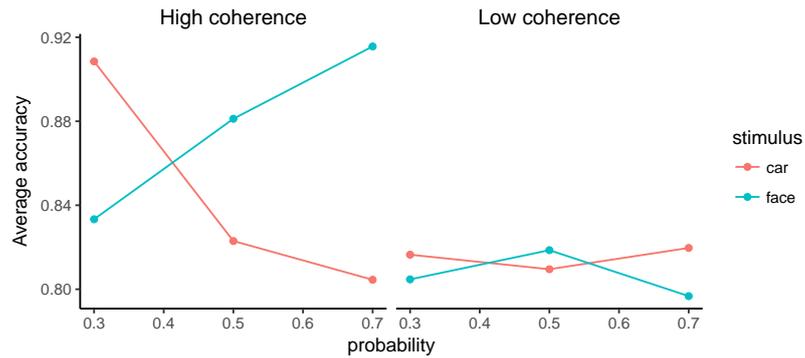


Figure B.2: Pilot experiment 1: Average accuracy of response and how it changes with stimulus type, cue probability, and stimulus coherence.

feedback after each trial. As before, this takes the form of a *correct* or *incorrect* symbol, which is displayed for 500ms. Pilots 1 and 3 took place in one session on one day. Pilot 2, on the other hand, took place in two sessions over consecutive days.

B.3 Pilot experiment 1

Pilot 1 used 6 participants (2 male, 4 female). Each participant was shown 6 blocks of trials, with each block containing 72 trials. The experimental task on each trial was exactly as described above. The two noise levels used were as follows: Low coherence for both face and car images, 32.50% coherence; high coherence for both face and car images, 37.50% coherence.

Accuracy analysis

The average accuracy of response and how it changes with predictor variables is shown in Figure B.2. Note that the three predictor variables here are a) the stimulus type on that trial, i.e., whether the stimulus is a face or car image, b) the coherence type on that trial, i.e., whether it is high or low coherence, and c) the probabilistic information provided by the cue on that trial. We represent the cue predictor variable as a continuous variable that takes on values of 0.30, 0.50 or 0.70. These values represent the probabilities of the upcoming stimulus on each trial being a face, and 1 minus the values represent that probability that the upcoming stimulus is a car.

We statistically model the accuracy of responses using a random effects (i.e., multi-level or hierarchical) logistic regression model. In particular, the logarithm of the odds of an accurate response on any given trial is modelled as a function of three fixed effects, and their interactions, and two random effects. The three fixed effect variables are the stimulus type on that trial, the coherence type on that trial, and the probabilistic information provided by the cue on that trial. The two random effects model a) the effect of the particular participant and b) the specific image that is shown on any given trial. The purpose of random effects for the participant is simply to model how each individual participant may have, on average, a higher or lower accuracy rate in this experiment. The random effect for the specific image models how a particular e.g. face image may have, on average, a higher or lower accuracy rate. To be clear, both the participant and image random effects are effectively random intercept terms in the regression, see below for more details. In other words, we do not model these effects as random slopes that could model the potential interaction effects of participant or image on the stimulus, coherence, or cue probability predictors.

We expect a potential three way interaction between the stimulus type, coherence level and cue probability variables. In other words, we expect that the probability of an accurate response will vary according to the cue probability and this will differ for trials with face stimuli and car stimuli, and also with the level of coherence level of the stimulus. Clearly, the effect of the cue is expected to affect accuracy of response differently depending on whether the trial has a face or a car stimulus. For example, if the cue is a triangle, then the probability of the upcoming stimulus being a car is 0.70 and the probability of it being a face is 0.30. If the actual stimulus on that trial is in fact a car, we expect a *more* accurate response, and we expect a *less* accurate if the actual stimulus is a face. As such, an interaction between the stimulus and cue is necessarily the case if the cue is in any way effective. In addition, we expect that this interaction may vary between the high and low coherence trials. This is not necessarily the case according to the main hypothesis under investigation. However, it is likely that the effect of cue on accuracy may be more or less effective depending on the level of noise in the stimulus.

Therefore, in this analysis, we obtain a measure of model fit for the full interaction model, i.e. the model that has the three-way interaction, the three two-way interactions,

and the three main effects. We compare this to the model fit measure of the model without the three-way interaction, i.e. with only the three two-way interactions, and the three main effects. Below, we refer to this model as the *two-way* model. For completeness, we also obtain the model fit measure for the model with only the main effects, referring to this as the *one-way* model, and also for the *null* model where the probability of an accurate response is modelled as a constant probability across all trials.

All three random effects logistic regression models can be represented using the following general form:

$$\log \left(\frac{P(y_i = 1)}{1 - P(y_i = 1)} \right) = \underbrace{\beta_0 + \beta_1 x_{1i} + \dots + \beta_K x_{Ki}}_{\text{fixed effects}} + \underbrace{\gamma_{[\text{subject}_i]}}_{\text{random effects}} + \underbrace{v_{[\text{item}_i]}}_{\text{random effects}} .$$

Here, the y_i signifies the accuracy of the response on trial i , with $y_i = 1$ when the response is correct and $y_i = 0$ when the response is incorrect. The $\beta_0, \beta_1 \dots \beta_K$ represent the coefficients for the fixed effects and their interactions. There is a variable number of these coefficients depending on which predictor variables and interactions are used in any given model. For each of the $j \in 1 \dots J$ participants in the experiment, there is γ_j , and $\gamma_1 \dots \gamma_j \dots \gamma_J$ are also modelled as drawn from a normal distribution with 0 mean and standard deviation σ_γ . Likewise, each of the $j = l \in 1 \dots L$ participants in the experiment, there is a v_l , and $v_1 \dots v_l \dots v_L$ are also modelled as drawn from a normal distribution with 0 mean and standard deviation σ_v .

In the following table, we show the AIC values for *Full*, *two-way*, *one-way*, and *null* models.

Full model	Two-way model	One-way model	Null model
2027	2038	2042	2059

As can be seen, the *Full* model has the best model fit. Each model fit can be compared the best fitting model using the differences in AIC values. Following Burnham, any model whose AIC value is greater than at least 10 units is said to have effectively zero evidence in favour of it relative to the best fitting model. As such, the *Full* model can be said to be *highly significantly* better than the other three models.

Given that there is a three-way interaction, the effect of the cue type on the probability of an accurate response is different for face and car stimuli, and this effect itself differs for the low and high coherence trials. In order to simplify the subsequent analysis, we now use two separate regression models, one for the low coherence trials and another for the high coherence trials. The table of coefficients for the predictors in the high coherence trials is as follows:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	3.48	0.53	6.60	<0.01
stimulusface	-2.16	0.60	-3.61	<0.01
probability	-2.91	0.76	-3.81	<0.01
stimulusface:probability	4.98	1.13	4.41	<0.01

As can be seen, there is a highly significant effect of stimulus type, cue probability and their interaction. The higher the probability, according to the cue, that the upcoming stimulus is a face, the higher the accuracy on the face trials. On the car trials, the higher the probability that the upcoming stimulus is a car (i.e. 1 minus the probability that it will be a face), the higher the accuracy of the response. These results are precisely in line with our hypothesis, and show that even with small sample sizes, the hypothesised effects in this experiment do in fact occur, at least on the high coherence trials.

On the low coherence trials, the effect of the cue and stimulus type and their interaction is as follows:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	1.58	0.45	3.50	<0.01
stimulusface	0.11	0.52	0.22	0.83
probability	0.26	0.70	0.37	0.71
stimulusface:probability	-0.38	0.98	-0.39	0.70

Here, there is no significant effect of cue, nor any interaction.

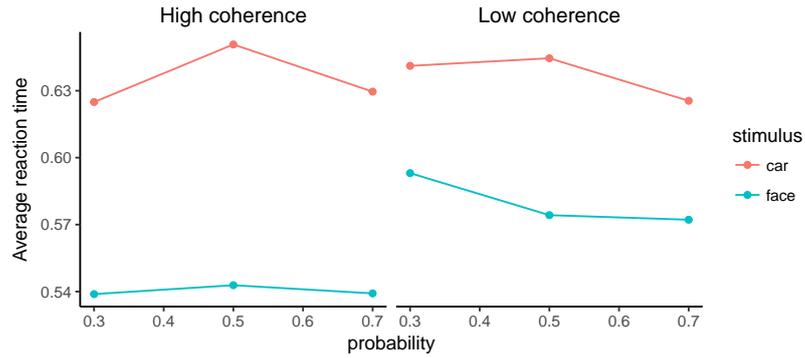


Figure B.3: Pilot experiment 1: Average reaction time and how it changes with stimulus type, cue probability, and stimulus coherence.

B.3.1 Reaction time analysis

The average reaction time on each trial and how it changes with predictor variables is shown in Figure B.3. The three predictor variables here are as above, i.e., the stimulus type on that trial, the coherence type on that trial, and the probabilistic information provided by the cue on that trial. Again, we represent the cue predictor variable as a continuous variable that represents the probability of the upcoming stimulus on each trial being a face.

In the Figure B.3, the results are shown for all trials, irrespective of whether the response was correct or inaccurate. In Figure B.4, we provide the same information as shown in Figure B.3 but where the trials are divided into accurate and inaccurate responses. In order to simplify the analysis, and given that there are accurate responses on 84% of the trials, we will now restrict our analysis of reaction times to just the accurate responses.

Following the approach to the accuracy analysis described above, we statistically model the reaction time of accurate responses using a random effects (i.e., multilevel or hierarchical) linear regression model. In particular, the logarithm of the reaction time on any given trial where the response is accurate is modelled as a function of three fixed effects, and their interactions, and two random effects. Just as in the case of the accuracy logistic random effects model, the three fixed effect variables are the stimulus type on that trial, the coherence type on that trial, and the probabilistic information provided by

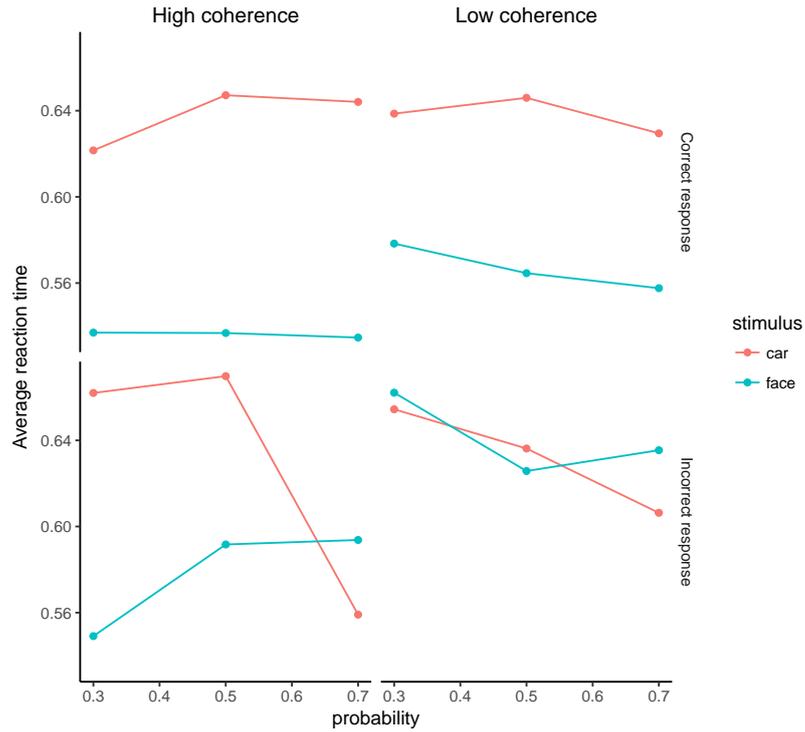


Figure B.4: Pilot experiment 1: Average reaction time and how it changes with stimulus type, cue probability, and stimulus coherence. Correct and incorrect responses are shown separately.

the cue on that trial, and the two random effects model are the effect of the particular participant and the specific image that is shown on any given trial. The rationale for the two random effects is identical to that in the case of the accuracy analysis model.

Again, we expect a potential three way interaction between the stimulus type, coherence level and cue probability variables. Therefore, in this analysis, as above, we obtain a measure of model fit for a full interaction model, i.e. the model that has the three-way interaction, the three two-way interactions, and the three main effects. We compare this to the model fit measure of the model without the three-way interaction, referred to by the *two-way* model, and the three main effects, referred to as the *one-way* model. We also use a null model whereby reaction time on accurate responses is treated as constant on all trials.

All three random effects linear regression models for the log of the reaction time on accurate trials can be represented using the following general form:

$$\log(y_i) = \underbrace{\beta_0 + \beta_1 x_{1i} + \dots + \beta_K x_{Ki}}_{\text{fixed effects}} + \underbrace{\gamma_{[\text{subject}_i]}}_{\text{random effects}} + \underbrace{v_{[\text{item}_i]}}_{\text{random effects}} + \epsilon_i.$$

Here, y_i signifies the reaction time on trial i . The $\beta_0, \beta_1 \dots \beta_K$ represent the coefficients for the fixed effects and their interactions. Just as in the case of the logistic regression for accuracy, there is a variable number of these coefficients depending on which predictor variables and interactions are used in any given model. For each of the $j \in 1 \dots J$ participants in the experiment, there is γ_j , and $\gamma_1 \dots \gamma_j \dots \gamma_J$ are also modelled as drawn from a normal distribution with 0 mean and standard deviation σ_γ . Likewise, each of the $j = l \in 1 \dots L$ participants in the experiment, there is a v_l , and $v_1 \dots v_l \dots v_L$ are also modelled as drawn from a normal distribution with 0 mean and standard deviation σ_v . Finally, the trial by trial error term is ε_i which is also modelled as normally distributed with a mean of 0 and standard deviation of σ .

The model fit measure that we will use here is again the AIC. In the following table, we show the AIC values for *Full*, *two-way*, *one-way*, and *null* models.

Full model	Two-way model	One-way model	Null model
-867.70	-868.80	-863.40	-599

As can be seen, the *two-way* model has the best model fit. Note that here, the AIC values are all negative, while in the accuracy analysis above, they are positive. This is inconsequential and we still interpret the meaning of the values in the same way, i.e., the lower the AIC, the better the model fit. Each other model's fit can be compared the best fitting model using the differences in AIC values. The *Full* model clearly does not have a better model fit than the *two-way*, and as such we can rule out the presence of a three-way interaction. The AIC of *one-way* model is greater than that of the *two-way* model. This indicates that the *two-way* model has a better fit than the *one-way* model. This difference is, according to the widely accepted thresholds detailed in Burnham, "significant". We can further analyse the difference between the *one-way* and the *two-way* model by a log-likelihood ratio test. This difference is also significant, $\chi^2_{[3]} = 11.33, p < 0.01$

As above, in order to simplify the subsequent analysis, we now use two separate regression models, one for the low coherence trials and another for the high coherence trials. The table of coefficients for the predictors in the high coherence trials is as follows:

Estimate	Std. Error	df	t value	Pr(> t)
----------	------------	----	---------	----------

(Intercept)	-0.52	0.04	14.25	-13.19	<0.01
stimulusface	-0.08	0.04	1127	-2.17	0.03
probability	0.09	0.05	1127	1.80	0.07
stimulusface:probability	-0.16	0.07	1127	-2.13	0.03

Turning our attention the low coherence trials, the coefficients for the fixed effects are as follows:

	Estimate	Std. Error	df	t value	Pr(> t)
(Intercept)	-0.46	0.04	13.84	-10.34	<0.01
stimulusface	-0.09	0.04	1044	-2.07	0.04
probability	-0.04	0.06	1044	-0.71	0.48
stimulusface:probability	-0.05	0.08	1044	-0.55	0.58

Here, there is no significant effect of cue, nor any interaction.

B.3.2 Discussion of Pilot Experiment 1

The hypothesis that the speed and accuracy of response will vary systematically with how well a pre-stimulus cue predicts a stimulus is partially supported by this pilot experiment. In particular, we clearly see an improvement in accuracy on relatively noise-less trials when the cue indicated that the stimulus had a higher probability. This was evident in both our stimulus types. This effect, however, was not observed at all on the low coherence trials. In terms of speed of response, on the relatively noise-less trials and when the responses were accurate, speed of response increased, but not significantly so, on those trials where the cue indicates that the stimulus had a higher probability. This effect was not clearly present on the low coherence trials. As such, in summary, we can conclude that the noise of the trials may make an important difference to the strength of this effect, and also that the effect may be stronger or more evident in terms of accuracy of response, rather than reaction time.

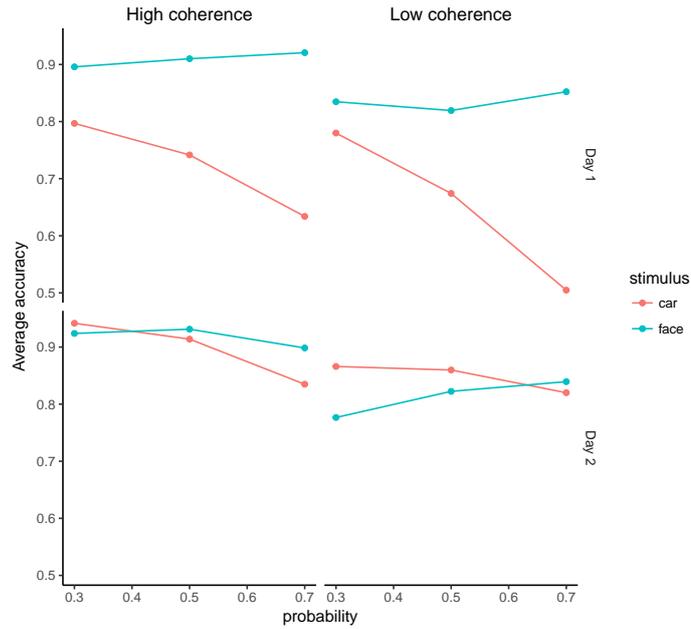


Figure B.5: Pilot experiment 2: Average accuracy of response and how it changes with stimulus type, cue probability, stimulus coherence, and training day.

B.4 Pilot experiment 2

The purpose of Pilot Experiment 2 was primarily to do a 2 consecutive day version of Pilot Experiment 1. In Pilot 1, as we have seen, only partial support for our main hypothesis was obtained. It was arguable, however, that the lack of the effects at relatively high noise levels and in the reaction time results was simply due to the fact insufficient time was given to learning the task. Our aim for the actual, i.e. non-pilot, experiment is to allow three days for learning, and we anticipate that the main phenomenon we are expecting to observe will become more evident with each day of learning. As such, it is reasonable to expect that the phenomena we partially observed in Pilot 1 will be clearer or more evident after two days of learning.

The only change that we made to the experiment was to change the luminance value of the Cathode Ray Tube (CRT) monitor.

B.4.1 Accuracy analysis

In order to analyze this data in a manner similar to the analysis used in experiment 1, we perform separate analyses for data from the first day and from the second day. In both cases, we perform a random effect logistic regression that is identical to the analysis used for the accuracy data in Pilot experiment 1. Given that the model is identical to that used in Pilot 1, we will not repeat its details here and instead refer the reader back to the details in the previous section.

In the following table, we show the AIC values for *Full*, *two-way*, *one-way*, and *null* models for the data on Day 1.

Full model	Two-way model	One-way model	Null model
1989	1987	2005	2119

While in the following table, we show the AIC values for *Full*, *two-way*, *one-way*, and *null* models for the data on Day 2.

Full model	Two-way model	One-way model	Null model
1448	1446	1453	1481

On both days, it is the *two-way* model that is the model with the best model fit. In other words, we can rule out the presence of the three way interaction between the stimulus type, cue probability and the coherence level. The *two-way* model is also “significantly” better than the *one-way* model too, given that the AIC value for the *two-way* model is lower than that of the *one-way* model by 17.25 units on Day 1 and by 6.89 units on Day 2. Note that here we follow Burnham and Anderson (2003) and interpret differences of greater than 4 units as indicating the superiority of the model with the lower value.

In order to further analyse the role of the three predictor variables, we begin by sequentially dropping each of the three pairwise interactions, i.e. the interactions of stimulus and cue probability, stimulus and coherence, probability and coherence, and assessing whether there is a rise in the AIC value. If there is a rise in the AIC with a particular pairwise interaction removed, this indicates that that pairwise interaction is improving the

model fit. On the other hand, if there is no change in the AIC value that indicates that the pairwise interaction is not improving the model fit and can be removed. In the following table, we see the AIC values of the model with all three pairwise interactions, and the model when each one of the three pairwise interactions is dropped.

Two-way model	Drop stim*prob	Drop stim*coh	Drop prob*coh
1987	2005	1988	1986

As can be seen, dropping the interaction between the stimulus and cue probability results in a substantially poorer model fit. Dropping the other two interactions, however, does not result in any substantial change in the AIC. In fact, dropping the interaction between the cue probability and coherence results in a small improvement in model fit. Note that an improvement in AIC is possible given that AIC penalises models with larger number of parameters. As such, even if the likelihood of model without a given interaction does not change, the AIC value may be lower due to there being less parameters in the model.

Dropping the two redundant interactions, the coefficients for the resulting model are as follows:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	2.45	0.30	8.13	<0.01
stimulusface	-0.53	0.41	-1.29	0.20
probability	-2.66	0.45	-5.85	<0.01
coherencelow	-0.45	0.11	-3.90	<0.01
stimulusface:probability	3.38	0.75	4.49	<0.01

The results for the model comparison procedure on the *two-way* model but for the data from the second day is as follows:

Two-way model	Drop stim*prob	Drop stim*coh	Drop prob*coh
1446	1449	1447	1451

Here, we see that when the interaction between coherence and cue probability is dropped, there a rise in the AIC value by about 5 units. When the stimulus and cue probability interaction is dropped, on the other hand, the rise is by approximately 3 units. This is a marginal result, given that we usually require differences in AIC values of 4 or more to be considered significant, while differences of less than 2 are not considered to be meaningful. Dropping the one redundant interaction, i.e. the stimulus and coherence interaction, from this model, the coefficients for the resulting model are as follows:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	4.13	0.53	7.77	<0.01
stimulusface	-1.22	0.48	-2.55	0.01
probability	-2.78	0.86	-3.22	<0.01
coherencelow	-1.78	0.49	-3.60	<0.01
coherencelow:probability	1.90	0.90	2.10	0.04
probability:stimulusface	2.19	0.89	2.45	0.01

On the basis of these accuracy results, it is clear that on Day 1, one of the main phenomena that we expect to observe, i.e. that accuracy increases whenever the cue has indicated that a face or car stimulus is more likely to occur, does in fact occur. There is a significant interaction between cue probability and stimulus type, with accuracy for face stimuli increasing whenever the cue indicates a higher probability of a face stimulus, and the accuracy for car stimuli increasing when the cue indicates a high probability of car stimulus. This phenomenon occurs at both the higher and lower noise conditions. However, as can be clearly seen in the figure above, the average accuracy for car stimuli is much lower than the accuracy for face stimuli, regardless of the cue probability. Moreover, this average accuracy is far lower than the accuracy observed for car stimuli during Pilot Experiment 1. We will return to this point in the discussion of Pilot 2 below.

In addition to the anomalous result concerning the accuracy of car stimuli, the fact that the role of cue probability on Day 2 is not as strong as that on Day 1 is not an encouraging result. In order to explore this result further, in Figure B.6, we have plotted the average accuracy on both days, as a function of the stimulus and cue probability and

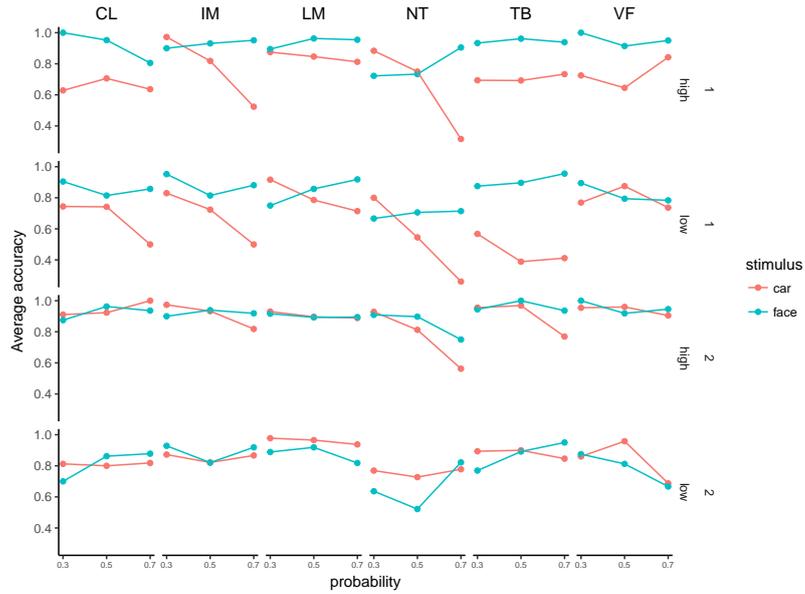


Figure B.6: Pilot experiment 2: Average accuracy of response and how it changes with training, stimulus type, cue probability, and stimulus coherence. Shown for each individual subject in the study.

collapsing over the coherence levels, individually for each subject. As can be seen in this figure, notwithstanding the considerable individual variability across participants, there appears to be marked improvement in accuracy by Day 2, particular on the car stimuli. It may be the case, therefore, that a ceiling effect is occurring by the second day of learning. This is particularly evident in the high coherence trials on Day 2.

B.4.2 Reaction time analysis

In the Figure B.7, we show the average reaction time as a function of cue probability, stimulus type, accuracy, and day of experiment.

Given that 84% of the trials were LM accurate, it what follows, we will restrict our attention to the data from just these trials. Likewise, following the nature of the analysis we performed on the accuracy data, we will perform two separate random effects regressions on the data from Day 1 and Day 2. On each day, we use an identical linear random effects regression model to analyse how the logarithm of reaction time varies with cue probability, stimulus, coherence, and their interactions. We use an identical linear random effects

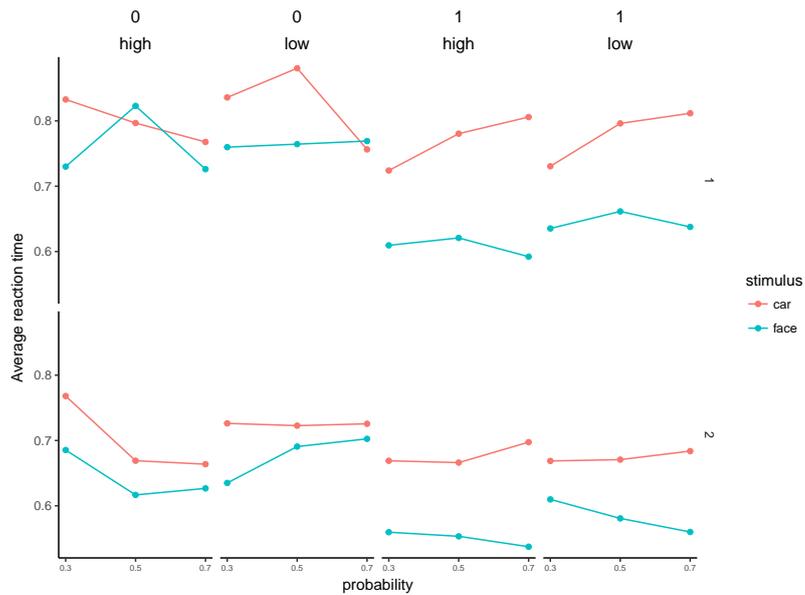


Figure B.7: Pilot experiment 2: Average reaction time and how it changes with training, stimulus type, cue probability, and stimulus coherence.

model to that used to analyse the reaction time in Pilot 1. As such, we will not repeat the details of the model here.

In the following table, using data from Day 1, we provide the AIC values for a set of four models, each with different types of interactions, exactly as we used in Pilot 1.

Full model	Two-way model	One-way model	Null model
-502.10	-503.60	-477.20	-146.80

In the following table, we provide the corresponding results using only the data from Day 2.

Full model	Two-way model	One-way model	Null model
-1346	-1348	-1328	-859.40

In both of these analyses, we see that the *two-way* model is the model with the lowest AIC value. Note, as was the case in the reaction time analysis for Pilot 1, the AIC values here are negative. This is inconsequential and does not affect our interpretation of the AIC values. As was done with the accuracy analysis, we can sequentially drop each of the three interaction terms from the *two-way* model.

The following table provides the AIC values for the model based on Day 1 data, where each one of the three pairwise interactions is dropped:

Two-way model	Drop stim*prob	Drop stim*coh	Drop prob*coh
-364.80	-349.20	-362.80	-366.60

As we can see, a substantial rise in the AIC value occurs when the interaction of stimulus and probability is dropped, but this does not occur for the other interactions. For the case of the data from Day 2, the AIC values for the models with each of the pairwise interactions dropped is as follows:

Two-way model	Drop stim*prob	Drop stim*coh	Drop prob*coh
-1284	-1276	-1275	-1286

We can see here that it is only when the interaction between cue probability and coherence is dropped is there little of no change on the AIC value of the model.

In order to analyse the reaction time data from Day 1, we will drop the interactions between stimulus and coherence and probability and coherence. The fixed effects coefficients from this model are as follows:

	Estimate	Std. Error	df	t value	Pr(> t)
(Intercept)	-0.39	0.04	17.99	-10.78	<0.01
stimulusface	-0.05	0.03	2002	-1.66	0.10
probability	0.17	0.05	2003	3.84	<0.01
coherencelow	0.04	0.01	1996	4.41	<0.01
stimulusface:probability	-0.27	0.06	2003	-4.23	<0.01

From this, we see a significant effect of coherence, cue probability, and interaction between stimulus and cue probability. Surprisingly, we do not obtain a significant main effect of stimulus type.

To analyse the reaction data from Day 2, we will drop the interaction between probability and coherence. The fixed effects coefficients from this model are as follows:

	Estimate	Std. Error	df	t value	Pr(> t)
(Intercept)	-0.43	0.05	8.78	-9.42	<0.01
stimulusface	-0.11	0.03	2127	-3.87	<0.01
probability	0.04	0.04	2129	1.08	0.28
coherencelow	<0.01	0.01	2124	0.44	0.66
stimulusface:probability	-0.16	0.05	2128	-3.26	<0.01
stimulusface:coherencelow	0.06	0.02	2124	3.81	<0.01

From this, we see a significant effect of stimulus type, and significant interactions between stimulus and cue probability and coherence and cue probability.

B.4.3 Conclusion from Pilot 2

In Pilot experiment 2, generally speaking, we observe the expected effect of cue probability on both accuracy and reaction time, at high and low coherence levels, and on both days of learning. Although this is generally encouraging, some results are surprising or anomalous. In particular, the accuracy level for car stimuli is markedly lower than for face stimuli on Day 1, but this accuracy level rises considerably by Day 2. In fact, by Day 2, we may be observing a ceiling effect. Given that Pilot 2 differs from Pilot 1 only in that Pilot 2 is a two day experiment, it is surprising that the accuracy level of cars in Pilot 2 Day 1 is very unlike that observed in Pilot 1. This may be due to the change in luminance of the monitor.

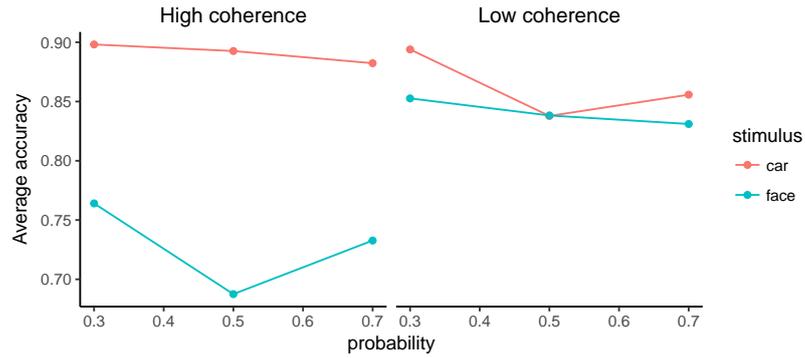


Figure B.8: Pilot experiment 3: Average accuracy of response and how it changes with stimulus type, cue probability, and stimulus coherence. Results shown separately for correct and incorrect responses.

B.5 Pilot experiment 3

In the previous pilot experiment, we observed a particularly low accuracy level for cars on Day 1, and then a near general ceiling effect of accuracy on Day 2. The primary purpose of Pilot experiment 3 is to assess whether higher noise for faces, but keeping the noise level of cars the same, would at least partially remedy this situation by making face stimuli more difficult to recognize. The noise levels for faces will be as follows: Low coherence, 27.50%; high coherence 32.50%. Our aim was to carry out this experiment on one day only, and see if we observe the expected cross-over interaction between cue probability and stimulus type. We also reset the luminance of the monitor to the original settings in Pilot 1. These two changes are the only changes we applied, and so Pilot 3 is identical to Pilot 1, and therefore its analysis will proceed in the same manner as the analysis in Pilots 1 and 2. As such, the details of the analysis are not repeated in full here (see above for more information).

B.5.1 Accuracy analysis

Figure B.8 shows the average accuracy as a function of the three predictors; stimulus type, coherence, and cue probability.

In the following table, we show the AIC values for *Full*, *two-way*, *one-way*, and *null* models (see above for definitions of these models).

Full model	Two-way model	One-way model	Null model
1468	1466	1476	1514

Here, we see that the *two-way* model is the preferred model according to its AIC value, and we can thus rule out the presence of three way interaction. As we did in Pilot 2 when we obtained a similar result, we further analyse the pairwise interactions in the *two-way* model.

Two-way model	Drop stim*prob	Drop stim*coh	Drop prob*coh
1466	1465	1478	1464

These results indicate that there is now an interaction between stimulus and cue probability, which is clearly not the expected result, nor is it in line with the results obtained in the previous pilot experiments. Dropping the redundant interactions, the coefficients for the predictors in the resulting model are as follows:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	2.56	0.41	6.29	<0.01
stimulusface	-0.73	0.15	-4.99	<0.01
probability	-0.94	0.59	-1.61	0.11
coherencelow	0.09	0.45	0.20	0.84
coherencelow:probability	0.52	0.83	0.62	0.53

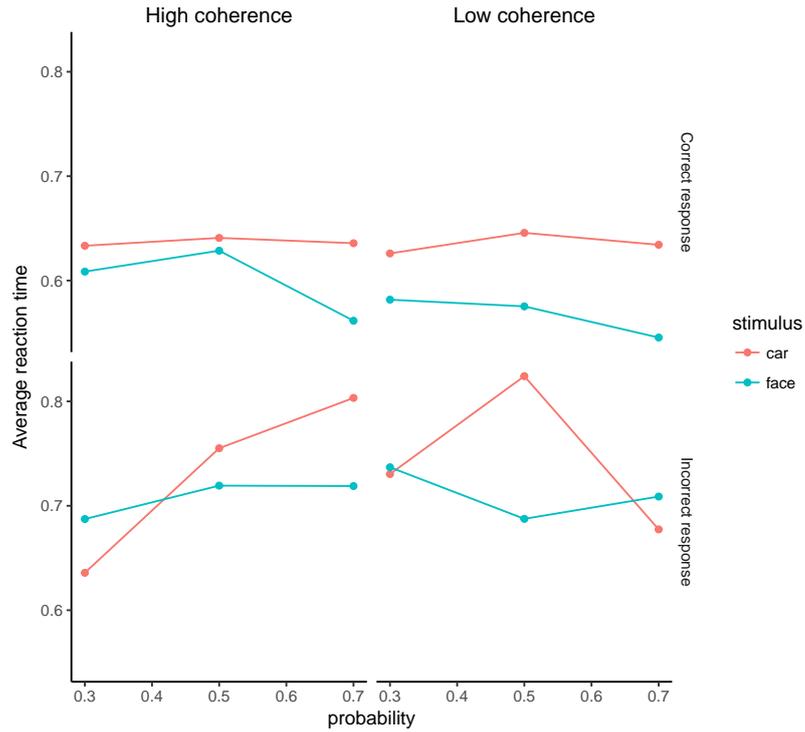


Figure B.9: Pilot experiment 3: Average reaction and how it changes with stimulus type, cue probability, and stimulus coherence. Results separated by accurate and inaccurate responses.

From this, we see that the only effect that is significant is the effect of the stimulus type. Neither the cue probability itself, nor coherence, nor their interaction was significant.

B.5.2 Reaction time analysis

Figure B.9 shows the average reaction time as a function of the three main predictors, for the accurate and inaccurate results.

As before, given that 83% of the trials were accurate, we will restrict the following analysis to only accurate responses. The following table provides the AIC values for our main models.

Full model	Two-way model	One-way model	Null model
-738.90	-740.80	-723.40	-613.70

Again, we see that the *two-way* model has the best AIC value. As before, we now analyse the roles the three pairwise interactions in the *two-way* model. Again, we show the AIC value when each one of the pairwise interactions is individually dropped.

As we can see from this, there is a marked rise in AIC when the stimulus and cue probability interaction is dropped, but a not substantial rise or even a decrease in AIC when the other two interaction are dropped. Dropping the redundant interactions results in a model with the following coefficients for the fixed effects:

Here, we see a significant interaction between stimulus type and cue probability, but no other effects are significant. The predicted reaction time according to this best model is as shown in the following figure.

B.5.3 Conclusion from Pilot experiment 3

The aim of Pilot 3 was primarily to assess whether changing the coherence level of face stimuli to lower values would remedy the surprising and anomalous results seen in Pilot 2. This was clearly not the case. In particular, the accuracy results show no sign of the expected effect of cue probability which have been observed, at least partially, in the accuracy results of the previous pilot experiments. Although these new coherence settings for faces may not necessarily eradicate the effect we are hoping to observe, they certainly do not make it clearer or enhanced.

B.6 General conclusions from pilot experiments

The three pilot experiments, taken as a whole, show that the effect of cue probability on either response speed or accuracy, or both, will be observed even with relatively low numbers of subjects, and in different noise conditions. Our main recommendations for the main experiments on the basis of the pilot experiments are therefore the following:

- We will use the 32.50% and 37.50% noise levels for both the face and car stimuli.
- We will not use any feedback in the main experiment. The possible ceiling effects we observe on Day 2 of learning may interfere with our analysis, and as such we will remove feedback so that learning is more challenging.

- The luminance setting of the monitor will be kept at the settings of Pilot 1 and 3.
- We will change the pre-stimulus cue from the three shapes used above to the statements 70F/30C, 50/50, 70C/30F, which denote 70% probability of face (30% probability of car), 50% probability of face (50% probability of face), and 70% probability of car (30% probability of face). We will explain to the subjects in the instructions what these statements mean. The purpose of making this change was to eliminate the inevitable learning process where the participants must learn the meaning of the cues.

References

- Aarts, E., Verhage, M., Veenliet, J. V., Dolan, C. V., & van der Sluis, S. (2014). A solution to dependency: using multilevel analysis to accommodate nested data. *Nature neuroscience*, *17*(4), 491–496.
- Ahissar, M., & Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature*, *387*(6631), 401–406.
- Ahissar, M., & Hochstein, S. (2004). The reverse hierarchy theory of visual perceptual learning. *Trends in cognitive sciences*, *8*(10), 457–464.
- Albright, T. D. (2012). On the perception of probable things: neural substrates of associative memory, imagery, and perception. *Neuron*, *74*(2), 227–245.
- Aslin, R. N., Saffran, J. R., & Newport, E. L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological science*, *9*(4), 321–324.
- Auckland, M. E., Cave, K. R., & Donnelly, N. (2007). Nontarget objects can influence perceptual processes during object recognition. *Psychonomic bulletin & review*, *14*(2), 332–337.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of memory and language*, *59*(4), 390–412.
- Baldassarre, A., Lewis, C. M., Committeri, G., Snyder, A. Z., Romani, G. L., & Corbetta, M. (2012). Individual variability in functional connectivity predicts performance of a perceptual task. *Proceedings of the National Academy of Sciences*, *109*(9), 3516–3521.
- Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. *Vision research*, *27*(6), 953–965.

- Bao, M., Yang, L., Rios, C., He, B., & Engel, S. A. (2010). Perceptual learning increases the strength of the earliest signals in visual cortex. *Journal of Neuroscience*, *30*(45), 15080–15084.
- Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, *5*(8), 617–629.
- Basso, M. A., & Wurtz, R. H. (1998). Modulation of neuronal activity in superior colliculus by changes in target probability. *Journal of Neuroscience*, *18*(18), 7519–7534.
- Basten, U., Biele, G., Heekeren, H. R., & Fiebach, C. J. (2010). How the brain integrates costs and benefits during decision making. *Proceedings of the National Academy of Sciences*, *107*(50), 21767–21772.
- Bates, D. M. (2010). lme4: Mixed-effects modeling with R. <http://lme4.r-forge.r-project.org/lmmwR/lrgprt.pdf>.
- Blank, H., Biele, G., Heekeren, H. R., & Philiastides, M. G. (2013). Temporal characteristics of the influence of punishment on perceptual decision making in the human brain. *The Journal of Neuroscience*, *33*(9), 3939–3952.
- Blanz, V., & Vetter, T. (1999). A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on computer graphics and interactive techniques* (pp. 187–194).
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological review*, *113*(4), 700.
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive psychology*, *57*(3), 153–178.
- Burnham, K. P., & Anderson, D. R. (2003). *Model selection and multimodel inference: a practical information-theoretic approach*. Springer Science & Business Media.
- Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, *14*(5), 365–376.
- Censor, N., Bonneh, Y., Arieli, A., & Sagi, D. (2009). Early-vision brain responses which predict human visual segmentation and learning. *Journal of Vision*, *9*(4), 12–12.

- Chen, N., Bi, T., Zhou, T., Li, S., Liu, Z., & Fang, F. (2015). Sharpened cortical tuning and enhanced cortico-cortical communication contribute to the long-term neural mechanisms of visual motion perceptual learning. *Neuroimage*, *115*, 17–29.
- Cravo, A. M., Rohenkohl, G., Wyart, V., & Nobre, A. C. (2013). Temporal expectation enhances contrast sensitivity by phase entrainment of low-frequency oscillations in visual cortex. *Journal of Neuroscience*, *33*(9), 4002–4010.
- Crist, R. E., Kapadia, M. K., Westheimer, G., & Gilbert, C. D. (1997). Perceptual learning of spatial localization: specificity for orientation, position, and context. *Journal of neurophysiology*, *78*(6), 2889–2894.
- Crist, R. E., Li, W., & Gilbert, C. D. (2001). Learning to see: experience and attention in primary visual cortex. *Nature neuroscience*, *4*(5), 519–525.
- Dakin, S. C., Hess, R. F., Ledgeway, T., & Achtman, R. L. (2002, Jul). What causes non-monotonic tuning of fMRI response to noisy images? *Current Biology*, *12*(14), R476–477.
- Davenport, J. L., & Potter, M. C. (2004). Scene consistency in object and background perception. *Psychological Science*, *15*(8), 559–564.
- de Lange, F. P., Rahnev, D. A., Donner, T. H., & Lau, H. (2013). Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *Journal of Neuroscience*, *33*(4), 1400–1410.
- Diaz, J. A., Queirazza, F., & Philiastides, M. G. (2017). Perceptual learning alters post-sensory processing in human decision making. *Nature Human Behaviour*, *1*, Article Number: 0035.
- Dosher, B. A., & Lu, Z.-L. (1999). Mechanisms of perceptual learning. *Vision research*, *39*(19), 3197–3221.
- Duda, R. O., Hart, P. E., & Stork, D. G. (2012). *Pattern classification*. John Wiley & Sons.
- Dunovan, K. E., Tremel, J. J., & Wheeler, M. E. (2014). Prior probability and feature predictability interactively bias perceptual decisions. *Neuropsychologia*, *61*, 210–221.

- Elbert, T., Pantev, C., Wienbruch, C., Rockstroh, B., & Taub, E. (1995). Increased cortical representation of the fingers of the left hand in string players. *Science*, *270*(5234), 305–307.
- Fahle, M. (2004). Perceptual learning: A case for early selection. *Journal of vision*, *4*(10), 4–4.
- Fahle, M. (2005). Perceptual learning: specificity versus generalization. *Current opinion in neurobiology*, *15*(2), 154–160.
- Fahle, M., & Edelman, S. (1993). Long-term learning in vernier acuity: Effects of stimulus orientation, range and of feedback. *Vision research*, *33*(3), 397–412.
- Fang, Y. (2011). Asymptotic equivalence between cross-validations and akaike information criteria in mixed-effects models. *Journal of Data Science*, *9*(1), 15–21.
- Farrell, S., & Lewandowsky, S. (2018). *Computational modeling of cognition and behavior*. Cambridge University Press.
- Fiorentini, A., & Berardi, N. (1980). Perceptual learning specific for orientation and spatial frequency. *Nature*, *287*(5777), 43–44.
- Fiser, J., & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, *99*(24), 15822–15826.
- Forstmann, B. U., Anwander, A., Schäfer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., ... Turner, R. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences*, *107*(36), 15916–15920.
- Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-sma facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences*, *105*(45), 17538–17542.
- Fouragnan, E., Retzler, C., Mullinger, K., & Philiastides, M. G. (2015). Two spatiotemporally distinct value systems shape reward-based learning in the human brain. *Nature communications*, *6*, 8107.
- Furmanski, C. S., Schluppeck, D., & Engel, S. A. (2004). Learning strengthens the response of primary visual cortex to simple patterns. *Current Biology*, *14*(7), 573–578.

- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models*. New York, NY: Cambridge University Press.
- Ghose, G. M., Yang, T., & Maunsell, J. H. (2002). Physiological correlates of perceptual learning in monkey v1 and v2. *Journal of Neurophysiology*, *87*(4), 1867–1888.
- Goldstone, R. L. (1998). Perceptual learning. *Annual review of psychology*, *49*(1), 585–612.
- Guggenmos, M., Wilbertz, G., Hebart, M. N., & Sterzer, P. (2016). Mesolimbic confidence signals guide perceptual learning in the absence of external feedback. *eLife*, *5*, e13388.
- Guitart-Masip, M., Huys, Q. J., Fuentemilla, L., Dayan, P., Duzel, E., & Dolan, R. J. (2012). Go and no-go learning in reward and punishment: interactions between affect and effect. *Neuroimage*, *62*(1), 154–166.
- Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., & Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *The Journal of Neuroscience*, *31*(17), 6339–6352.
- Heitz, R. P. (2014). The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Frontiers in neuroscience*, *8*, 150.
- Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, *36*(5), 791–804.
- Huys, Q. J., Guitart-Masip, M., Dolan, R. J., & Dayan, P. (2015). Decision-theoretic psychiatry. *Clinical Psychological Science*, *3*(3), 400–421.
- Huys, Q. J., Maia, T. V., & Frank, M. J. (2016). Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature neuroscience*, *19*(3), 404.
- Hyman, S. E., Malenka, R. C., & Nestler, E. J. (2006). Neural mechanisms of addiction: The role of reward-related learning and memory. *Annual Review of Neuroscience*, *29*, 565-598.
- Ivanoff, J., Branning, P., & Marois, R. (2008). fmri evidence for a dual process account of the speed-accuracy tradeoff in decision-making. *PLoS ONE*, *3*(7), e2635.
- Jehee, J. F., Ling, S., Swisher, J. D., van Bergen, R. S., & Tong, F. (2012). Perceptual learning selectively refines orientation representations in early visual cortex. *The Journal of neuroscience*, *32*(47), 16747–16753.

- Kahnt, T., Grueschow, M., Speck, O., & Haynes, J.-D. (2011). Perceptual learning and decision-making in human medial frontal cortex. *Neuron*, *70*(3), 549–559.
- Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences*, *88*(11), 4966–4970.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of American Statistical Association*, *90*(430), 773–795.
- Kelly, S. P., & O’Connell, R. G. (2013). Internal and external influences on the rate of sensory evidence accumulation in the human brain. *The Journal of Neuroscience*, *33*(50), 19434–19441.
- Kuai, S.-G., Levi, D., & Kourtzi, Z. (2013). Learning optimizes decision templates in the human visual cortex. *Current Biology*, *23*(18), 1799–1804.
- Lagerlund, T. D., Sharbrough, F. W., & Busacker, N. E. (1997). Spatial filtering of multi-channel electroencephalographic recordings through principal component analysis by singular value decomposition. *Journal of clinical neurophysiology*, *14*(1), 73–82.
- Law, C.-T., & Gold, J. I. (2008a). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature neuroscience*, *11*(4), 505–513.
- Law, C.-T., & Gold, J. I. (2008b). Neural Correlates of Perceptual Learning in a Sensory-Motor but not a Sensory Cortical Area. *Nature Neuroscience*, *11*(4), 505–513.
- Law, C.-T., & Gold, J. I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nature neuroscience*, *12*(5), 655–663.
- Lewis, C. M., Baldassarre, A., Committeri, G., Romani, G. L., & Corbetta, M. (2009). Learning sculpts the spontaneous activity of the resting human brain. *Proceedings of the National Academy of Sciences*, *106*(41), 17558–17563.
- Li, S., Mayhew, S. D., & Kourtzi, Z. (2009). Learning shapes the representation of behavioral choice in the human brain. *Neuron*, *62*(3), 441–452.
- Li, W., Piëch, V., & Gilbert, C. D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nature neuroscience*, *7*(6), 651.

- Lou, B., Li, Y., Philiastides, M. G., & Sajda, P. (2014). Prestimulus alpha power predicts fidelity of sensory encoding in perceptual decision making. *Neuroimage*, *87*, 242–251.
- Lu, Z.-L., Liu, J., & Doshier, B. A. (2010). Modeling mechanisms of perceptual learning with augmented hebbian re-weighting. *Vision research*, *50*(4), 375–390.
- Molchan, S. E., Sunderland, T., McIntosh, A., Herscovitch, P., & Schreurs, B. G. (1994). A functional anatomical study of associative learning in humans. *Proceedings of the National Academy of Sciences*, *91*(17), 8122–8126.
- Newsome, W. T., Britten, K. H., & Movshon, J. A. (1989). Neuronal correlates of a perceptual decision. *Nature*, *341*(6237), 52–54.
- O’Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012, Dec). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, *15*(12), 1729–1735.
- O’Doherty, J., 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.
- O’Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*, *304*(5669), 452–454.
- Oliva, A., & Torralba, A. (2007). The role of context in object recognition. *Trends in cognitive sciences*, *11*(12), 520–527.
- Palmer, t. E. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition*, *3*(5), 519–526.
- Pantev, C., Oostenveld, R., Engelien, A., Ross, B., Roberts, L. E., & Hoke, M. (1998). Increased auditory cortical representation in musicians. *Nature*, *392*(6678), 811–814.
- Parra, L., Alvino, C., Tang, A., Pearlmutter, B., Yeung, N., Osman, A., & Sajda, P. (2002). Linear spatial integration for single-trial detection in encephalography. *Neuroimage*, *17*(1), 223–230.

- Parra, L., Alvino, C., Tang, A., Pearlmutter, B., Yeung, N., Osman, A., & Sajda, P. (2003). Single-trial detection in eeg and meg: Keeping it linear. *Neurocomputing*, *52*, 177–183.
- Parra, L., Spence, C. D., Gerson, A. D., & Sajda, P. (2005). Recipes for the linear analysis of eeg. *Neuroimage*, *28*(2), 326–341.
- Peirce, J. W. (2007). Psychopy - psychophysics software in python. *Journal of neuroscience methods*, *162*(1), 8–13.
- Perna, A., Tosetti, M., Montanaro, D., & Morrone, M. C. (2008). Bold response to spatial phase congruency in human brain. *Journal of Vision*, *8*(10), 15–15.
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*(7106), 1042–1045.
- Petrov, A. A., Doshier, B. A., & Lu, Z.-L. (2005). The dynamics of perceptual learning: an incremental reweighting model. *Psychological review*, *112*(4), 715.
- Petrov, A. A., Van Horn, N. M., & Ratcliff, R. (2011). Dissociable perceptual-learning mechanisms revealed by diffusion-model analysis. *Psychonomic bulletin & review*, *18*(3), 490–497.
- Philiastides, M. G., Biele, G., Vavatzanidis, N., Kazzer, P., & Heekeren, H. R. (2010). Temporal dynamics of prediction error processing during reward-based decision making. *Neuroimage*, *53*(1), 221–232.
- Philiastides, M. G., Diaz, J. A., & Gherman, S. (2015). Spatiotemporal characteristics and modulators of perceptual decision-making in the human brain. In J.-C. Dreher & L. Tremblay (Eds.), *Decision neuroscience - handbook of reward and decision making*. Academic Press.
- Philiastides, M. G., Heekeren, H. R., & Sajda, P. (2014). Human scalp potentials reflect a mixture of decision-related signals during perceptual choices. *Journal of Neuroscience*, *34*(50), 16877–16889.
- Philiastides, M. G., Ratcliff, R., & Sajda, P. (2006a). Neural Representation of Task Difficulty and Decision Making During Perceptual Categorization: A Timing Diagram. *Journal Of Neuroscience*, *26*(35), 8965-8975.

- Philiastides, M. G., Ratcliff, R., & Sajda, P. (2006b). Neural representation of task difficulty and decision making during perceptual categorization: a timing diagram. *The Journal of Neuroscience*, *26*(35), 8965–8975.
- Philiastides, M. G., & Sajda, P. (2006). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral Cortex*, *16*(4), 509–518.
- Philiastides, M. G., & Sajda, P. (2007). Eeg-informed fmri reveals spatiotemporal characteristics of perceptual decision making. *The Journal of Neuroscience*, *27*(48), 13082–13091.
- Poggio, T., Fahle, M., & Edelman, S. (1992). Fast perceptual learning in visual hyperacuity. *Science*, *256*(5059), 1018–1021.
- Pourtois, G., Rauss, K. S., Vuilleumier, P., & Schwartz, S. (2008). Effects of perceptual learning on primary visual cortex activity in humans. *Vision research*, *48*(1), 55–62.
- Puri, A. M., Wojciulik, E., & Ranganath, C. (2009). Category expectation modulates baseline and stimulus-evoked activity in human inferotemporal cortex. *Brain research*, *1301*, 89–99.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, *85*(2), 59.
- Ratcliff, R., Gomez, P., & McKoon, G. (2004). A diffusion model account of the lexical decision task. *Psychological review*, *111*(1), 159.
- Ratcliff, R., Love, J., Thompson, C. A., & Opfer, J. E. (2012). Children are not like older adults: A diffusion model analysis of developmental changes in speeded responses. *Child development*, *83*(1), 367–381.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural computation*, *20*(4), 873–922.
- Ratcliff, R., Philiastides, M. G., & Sajda, P. (2009a). Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the eeg. *Proceedings of the National Academy of Sciences*, *106*(16), 6539–6544.
- Ratcliff, R., Philiastides, M. G., & Sajda, P. (2009b, APR 21). Quality of Evidence for Perceptual Decision Making is Indexed by Trial-to-Trial Variability of the EEG.

- Proceedings of the National Academy of Sciences of the United States of America*, 106(16), 6539-6544.
- Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological review*, 111(2), 333.
- Ratcliff, R., & Smith, P. L. (2010). Perceptual discrimination in static and dynamic noise: the temporal relation between perceptual encoding and decision making. *Journal of Experimental Psychology: General*, 139(1), 70.
- Ratcliff, R., Smith, P. L., & McKoon, G. (2015). Modeling regularities in response time and accuracy data with the diffusion model. *Current directions in psychological science*, 24(6), 458–470.
- Ratcliff, R., & Van Dongen, H. P. (2011). Diffusion model for one-choice reaction-time tasks and the cognitive effects of sleep deprivation. *Proceedings of the National Academy of Sciences*, 108(27), 11285–11290.
- Recanzone, G. H., Merzenich, M. M., Jenkins, W. M., Grajski, K. A., & Dinse, H. R. (1992). Topographic reorganization of the hand representation in cortical area 3b owl monkeys trained in a frequency-discrimination task. *Journal of Neurophysiology*, 67(5), 1031–1056.
- Recanzone, G. H., Schreiner, C. E., & Merzenich, M. M. (1993). Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. *The Journal of Neuroscience*, 13(1), 87–103.
- Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of neuroscience*, 22(21), 9475–9489.
- Rushworth, M. F., Mars, R. B., & Summerfield, C. (2009). General mechanisms for making decisions? *Current opinion in neurobiology*, 19(1), 75–83.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 1926–1928.
- Sagi, D., & Tanne, D. (1994). Perceptual learning: learning to see. *Current opinion in neurobiology*, 4(2), 195–199.
- Sasaki, Y., & Watanabe, T. (2017). When perceptual learning occurs. *Nature Human Behaviour*, 1, Article number: 0048 (2017).

- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80(1), 1-27.
- Schultz, W., Dayan, P., & Montague, P. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593-1599.
- Schwartz, S., Maquet, P., & Frith, C. (2002). Neural correlates of perceptual learning: a functional mri study of visual texture discrimination. *Proceedings of the National Academy of Sciences*, 99(26), 17137–17142.
- Seitz, R. J., & Roland, P. E. (1992). Learning of sequential finger movements in man: a combined kinematic and positron emission tomography (pet) study. *European Journal of Neuroscience*, 4(2), 154–165.
- Shibata, K., Sagi, D., & Watanabe, T. (2014). Two-stage model in perceptual learning: toward a unified theory. *Annals of the New York Academy of Sciences*, 1316(1), 18–28.
- Shima, K., & Tanji, J. (1998). Both supplementary and presupplementary motor areas are crucial for the temporal organization of multiple movements. *Journal of neurophysiology*, 80(6), 3247–3260.
- Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in neurosciences*, 27(3), 161–168.
- Strayer, D. L., & Kramer, A. F. (1994). Strategies and automaticity: I. basic findings and conceptual framework. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(2), 318.
- Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: neural and computational mechanisms. *Nature Reviews Neuroscience*.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT Press. Cambridge, MA.
- Troje, N. F., & Bühlhoff, H. H. (1996). Face recognition under varying poses: The role of texture and shape. *Vision research*, 36(12), 1761–1771.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological review*, 108(3), 550.
- Vandekerckhove, J., Tuerlinckx, F., & Lee, M. D. (2011). Hierarchical diffusion models for two-choice response times. *Psychological methods*, 16(1), 44.

- van Maanen, L., Brown, S. D., Eichele, T., Wagenmakers, E.-J., Ho, T., Serences, J., & Forstmann, B. U. (2011). Neural correlates of trial-to-trial fluctuations in response caution. *The Journal of Neuroscience*, *31*(48), 17488–17495.
- Watanabe, S. (2010). Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, *11*(Dec), 3571–3594.
- Watanabe, T., & Sasaki, Y. (2015). Perceptual learning: toward a comprehensive theory. *Annual review of psychology*, *66*, 197.
- Wiecki, T., Sofer, I., & Frank, M. (2013). Hddm: Hierarchical bayesian estimation of the drift-diffusion model in python. *Frontiers in Neuroinformatics*, *7*, 14.
- Wise, R., & Rompre, P. (1989). Brain Dopamine and reward. *Annual Review of Psychology*, *40*, 191-225.
- Xiao, L.-Q., Zhang, J.-Y., Wang, R., Klein, S. A., Levi, D. M., & Yu, C. (2008). Complete transfer of perceptual learning across retinal locations enabled by double training. *Current Biology*, *18*(24), 1922–1926.
- Yan, Y., Rasch, M. J., Chen, M., Xiang, X., Huang, M., Wu, S., & Li, W. (2014). Perceptual training continuously refines neuronal population codes in primary visual cortex. *Nature neuroscience*, *17*(10), 1380–1387.
- Yang, T., & Maunsell, J. H. (2004). The effect of perceptual learning on neuronal responses in monkey visual area v4. *The journal of Neuroscience*, *24*(7), 1617–1626.
- Zhang, J., & Rowe, J. B. (2014). Dissociable mechanisms of speed-accuracy tradeoff during visual perceptual learning are revealed by a hierarchical drift-diffusion model. *Frontiers in neuroscience*, *8*.