

BIROn - Birkbeck Institutional Research Online

Enabling Open Access to Birkbeck's Research Degree output

The perceptual prediction paradox : do actions exert a distinct influence on perception?

<https://eprints.bbk.ac.uk/id/eprint/47624/>

Version: Full Version

Citation: Thomas, Emily R. (2021) The perceptual prediction paradox : do actions exert a distinct influence on perception? [Thesis] (Unpublished)

© 2020 The Author(s)

All material available through BIROn is protected by intellectual property law, including copyright law.

Any use made of the contents should comply with the relevant law.

[Deposit Guide](#)
Contact: [email](#)

**The Perceptual Prediction Paradox: Do
Actions Exert a Distinct Influence on
Perception?**

Emily R. Thomas

Birkbeck, University of London

Thesis submitted to the Department of Psychological
Sciences for the degree of *Doctor of Philosophy*

March 2021

Declaration

I, Emily R. Thomas, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis through appropriate citation.

Experiment 6 reported in this thesis describes the conjoint work of a fellow post-doctoral lab member and myself. The project was an fMRI experiment where we both played an equivalent role in design, data collection and pre-processing, but the later analysis steps were conducted by the post-doctoral researcher. I was equally involved in interpretation and we will be writing the manuscript together.

March 2021

Peer-reviewed work

The work presented in Chapters 2 and 3 (excluding Experiment 4) of this thesis has been accepted for publication:

Thomas, E.R., Yon, D., Lange, F.P., Press, C. (in press). Action enhances predicted touch. *Psychological Science*.

Acknowledgements

I would like to thank the following individuals for their help throughout my PhD, which I could not have successfully completed without their guidance, support and friendship.

Firstly, I would like to say a big thank you to my supervisor Clare Press for your direction and wisdom over the past three and a half years. I am incredibly grateful for all your guidance, encouragement and good humour throughout. Thank you to Daniel Yon, Carl Bunce, James Chard and others in the Action Lab for your shared expertise, thought-provoking discussions and fun lab socials. Also to Jess Nicholson for all your help the past year, very appreciated. I would especially like to extend thanks to collaborators Joost Haarmsa and Peter Kok for all your guidance, direction and wisdom on the 7T project. Additionally to Matthew Longo, Floris de Lange, Sam Gilbert and staff at the WCHN and Birkbeck, your help throughout has been much appreciated.

I am very grateful for my network of friends and colleagues who have provided unconditional support through the highs and lows, which I could not have done this PhD without. A special thank you to Brittney for always being there for me, supporting me and keeping me sane throughout the last year of lock-downs. Kathryn and Gurmukh, thank you for being the best desk mates and for your encouragement as I approached the final stretch. Thanks also to Jenny and Maria for all your help and guidance, Susanne, Ellie, Aude, Louisa, Jess, Allison and the rest of the BMA family for our interesting discussions, your support, the laughs, and most of all for all the fun times.

Finally, I would like to thank my friends and family for all your support over the years. My biggest appreciation of all goes to Owen, for fuelling me with endless baked goods and for always asking me curious questions about tickling, but primarily for your undying encouragement and support, without which I would have lost sight of the finish line long ago.

Abstract

Perception of expected action outcomes has been thought for decades to be attenuated or ‘cancelled’. Underlying mechanisms are proposed to highlight informative unexpected events that signal the need to take new courses of action and update our models of the world. These accounts are thought to explain why we cannot tickle ourselves and unusual aspects of action in clinical populations. However, theories outside of action purport that predicted events are perceptually enhanced, allowing us to generate largely accurate representations of our noisy sensory world. This thesis presents a series of experiments investigating whether action prediction mechanisms necessarily generate attenuation, or whether they can enhance perception – as typically observed in domains outside of action. Chapter 1 outlines these theoretical accounts and evaluates the range of supporting literature across domains.

Chapters 2 and 3 examine how action predictions influence tactile perception. Chapter 2 uses a force-judgement paradigm that typically demonstrates attenuation and controls for some potential confounds. Chapter 3 explicitly dissociates the predictive contribution of these effects with manipulations of statistical probability from broader sensory cognition domains. Finding in these experiments that action predictions can enhance, rather than attenuate tactile perception, Chapter 4 sought to investigate the neural mechanisms underlying visuomotor prediction in Primary Visual Cortex using functional magnetic resonance imaging (fMRI). Univariate and multivariate analyses demonstrated that expected outcomes are represented with superior fidelity than unexpected outcomes. Chapter 5 used high-resolution (7T) fMRI to investigate whether a key assumption of predictive coding theories could reconcile the opposing neural prediction effects reported in the literature.

Chapter 6 summarises the findings of this thesis that sensorimotor predictions formed on the basis of statistical learning do not appear to exhibit a distinct influence on perception from other forms of prediction, and are instead likely generated by comparable underlying mechanisms.

Table of Contents

Chapter 1: Introduction	16
1.1. What are ‘top-down’ effects and how is ‘prediction’ defined?	17
1.2. How do predictions influence perception?.....	20
1.2.1. Predictions enhance perception.....	20
1.2.2. Predicted action outcomes are perceptually attenuated	22
1.2.3. Do action predictions influence perception in a distinct fashion?	25
1.2.4. Do action predictions influence tactile perception in a special way?	28
1.2.5. Open questions	29
1.3. The predictive nature of cortical effects – does the perceptual paradox extend to neural mechanisms?	30
1.3.1. Suppressed neural response to predicted consequences of action	30
1.3.2. Sharpening of expected sensory representations.....	32
1.3.3. Open questions	36
1.4. Could distinctly operating prediction mechanisms across cortical layers reconcile the neural paradox?	36
1.4.1. Hierarchical theories of prediction in the cortical column.....	36
1.4.2. Evidence for distinct prediction and prediction error populations.....	39
1.4.3. Open questions	41
1.5. Thesis outline	42
Chapter 2: Examining tactile attenuation in action.....	44
2.1. Experiment 1	45
2.1.1. Method	49
2.1.2. Results and Discussion.....	54
2.2. General Discussion.....	54
Chapter 3: Does expectation enhance or attenuate tactile perception?	59
3.1. Experiment 2	61
3.1.1. Methods.....	62

3.1.2. Results and Discussion.....	65
3.2. Experiment 3	65
3.2.1. Methods.....	66
3.2.2. Results and Discussion.....	68
3.3. Computational modelling.....	70
3.3.1. Methods.....	71
3.3.2. Results and Discussion.....	72
3.4. Experiment 4	76
3.4.1. Methods.....	78
3.4.2. Results and Discussion.....	79
3.5. General Discussion.....	82
Chapter 4: The nature of neural prediction mechanisms in action	88
4.1. Experiment 5	92
4.1.1. Methods.....	97
4.1.2. Results and Discussion.....	100
4.2. Experiment 6	105
4.2.1. Methods and Analyses	107
4.2.2. Results and Discussion.....	113
4.3. General Discussion.....	117
Chapter 5: The laminar profile of surprise	121
5.1. Experiment 7	125
5.1.1. Methods.....	127
5.1.2. Results and Discussion.....	130
5.2. Experiment 8	135
5.2.1. Methods and Analyses	141
5.2.2. Results and Discussion.....	150
5.3. General Discussion.....	157

Chapter 6: General Discussion	162
6.1. Thesis summary	162
6.2. How could the perceptual prediction paradox be reconciled?	166
6.2.1. Sensory gating.....	166
6.2.2. Repetition suppression	170
6.2.3. Expectation and signal strength	173
6.3. Expectation vs Attention.....	176
6.4. Implications for wider applications of Cancellation theories in action.....	177
6.5. Conclusions.....	178
References	179

List of Figures

Figure 1.1. Left: Hollow-face illusion, concave faces are perceived as convex when full shading cues are available (e.g., front on, left), adapted from Press et al. (2020b).

Middle: the Necker cube can be perceived in two different 3D perspectives. Right: after ‘explaining away’ the leaves in ambiguous input the tiger stands out more (taken from Kok & De Lange, 2015)..... 17

Figure 1.2. (A) Cancellation theories propose that neural activity associated with expected outcomes of action is suppressed in order to highlight unexpected events, resulting in expectation suppression effects that are greatest in populations tuned towards the presented stimuli. (B) Sharpening theories propose that we ‘pre-activate’ populations tuned to the expected stimulus and relatively suppress processing in populations representing other events, resulting in a sharpened neural response (adapted from Press et al., 2020b)..... 33

Figure 1.3. (A) Depiction of the six cortical layers intersected onto a section of occipital cortex. (B) Arrows denote the boundary between V1 and V2 (Secondary Visual Cortex). Adapted from Wagstyl et al. (2020). (C) Schematic representation of proposed extrinsic feedforward (red) and feedback (blue) connections across layers, between the LGN, V1 and V2; taken from Lawrence, Formisano, et al. (2019)..... 38

Figure 2.1. (A) On each trial, participants made downward movements with their right index finger either over a motion tracker (No Contact condition), or towards a button with which they made contact (Contact condition). Each movement elicited a tactile punctate event to the left index finger positioned directly below. (B) PSEs were calculated for each participant and this data represents an example participant in the No Contact condition for Active [dark blue] and Passive [light blue] trials. (C) Mean PSEs (\pm SEM) were higher in Active than Passive trials in the Contact conditions, but lower in Active than Passive trials in the No Contact condition. Lower PSEs indicate more intense target percepts (* $p < .05$, ** $p = .001$). (D) PSE effect of movement (Passive – Active) for the Contact (top) and No Contact (bottom) condition, plotted with raincloud plots displaying probability density estimates (upper) and box and scatter plots (lower). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive effects of movement indicate more intensely perceived active target events relative to passive

events, but negative values indicate the reverse – less intensely perceived active events..... 53

Figure 3.1. (A) Experiment 2 set-up. Participants made a downwards or upwards movement with their right index finger over a motion tracker, eliciting punctate tactile events to the left index or middle finger. (B) Experiment 3 set-up. Participants made only downwards movements, now with either their right index or middle finger. (C) Example training procedure for Experiment 2. Movements were perfectly predictive of tactile events in the training session, and 66.7% predictive in the test session. The same procedure was adopted in Experiment 3 but with different action types. Flash symbols illustrate location of stimulation. (D) Mean PSEs (\pm SEM) were lower for expected than unexpected trials in both Experiment 2 and Experiment 3, indicating more intensely perceived target stimuli (* $p < .05$). (E) PSE expectation effect (Unexpected – Expected) plotted with raincloud plots (Allen et al., 2019) displaying probability density estimates (upper) and box and scatter plots (lower), for Experiment 2 (top) and Experiment 3 (bottom). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive effect values indicate more intensely perceived expected events..... 64

Figure 3.2. (A) In an unbiased decision process (black lines), sensory evidence is accumulated towards the upper response boundary when stimuli are stronger than average (solid lines) and towards the lower response boundary when weaker than average (dotted lines). The starting point of evidence accumulation could however be shifted by decision circuits towards the upper boundary for expected events from the outset of the response window (influencing the parameter z ; blue lines - Start bias model). (B) Alternatively, altering the weights on sensory channels could bias evidence accumulation in line with expectations (influencing parameter db ; red lines – Drift bias model). Adapted from Yon et al., (2020). (C) Simulated Start + Drift bias (winning DIC model) expectation effect plotted against the empirical expectation effect, showing a significant positive correlation. (D) Simulated Drift bias expectation effect plotted against the empirical expectation effect accounting for simulated Start bias effects (plotted as the residuals from a model where the simulated Start bias effect predicts the empirical effect), again showing a significant positive correlation. Importantly, the regression analysis revealed that drift biases accounted for significant additional

variance once accounting for start biases. All expectation effects were calculated by subtracting Expected PSEs from the Unexpected PSEs..... 75

Fig 3.3. (A) Mean PSEs (\pm SEM) were lower for expected than unexpected trials for short vibration durations but not different for long duration vibrations (** $p < .01$). Lower PSEs indicate more intensely perceived tactile events. (B) PSE expectation effects (Unexpected – Expected) plotted with raincloud plots (Allen et al., 2019) displaying probability density estimates (upper) and box and scatter plots (lower), for short (top) and long (bottom) duration conditions. Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive expectation effect values indicate more intensely perceived expected events relative to unexpected events..... 81

Figure 4.1. (A) Cancellation theories propose that expected action–outcome representations are suppressed in order to highlight processing of unexpected events. This expectation suppression will be largest for units tuned towards the expected stimulus. (B) Sharpening theories propose that we activate expected representations and relatively suppress processing in units representing other features or events. (C) These processes will result in reduced contrast in activity patterns for expected outcomes relative to unexpected outcomes in the former (cancellation), yet increased contrast in expected activity patterns in the latter to reflect sharpened sensory representations (e.g., high signal-to-noise ratio; adapted from Press et al., 2020)..... 91

Figure 4.2. (A) A centrally presented visual cue instructed participants to abduct either their index or little finger which produced a CW or CCW grating stimulus. Participants could be asked either whether the stimulus was oriented CW or CCW in part one of the test trials, or whether it was oriented nearer to the blue or orange boundary lines in part two of the test trials. (B) One of two actions was executed on any given trial, whereby an index or little finger abduction perfectly predicted the orientation (CW or CCW) of each stimulus during training. The first part of the test session proceeded with an identical trial sequence except that actions now predicted stimulus orientations with 85% validity. Actions in part 2 of the test session predicted stimulus orientations with 50% validity. (C) In the second part of the test session participants observed the same trial sequence however were instead asked whether the stimulus orientation was nearer to the blue (63° or 117°) or orange (45° or 135°) boundary line, presented every mini-block. (D) A schematic for expected peak shift effects modelled using psychometric

functions. A bias towards perceiving extreme exemplars of the expected category (e.g., CW stimuli) would lead to a shift of the psychometric curve to the left (turquoise line), reflecting higher probability of responses that stimulus orientations were nearer to the more horizontal boundary (e.g., orange reference line, 45°)..... 96

Figure 4.3. (A) Left panel: Mean PSEs (\pm SEM) for expected and unexpected trials during the second part of the test session, which were not significantly different. PSEs were calculated on the probability of responses that stimulus orientations were ‘more horizontal’. PSE y-axis units index stimulus orientation, ranging from the most vertical (stimulus level 1) to the most horizontal (stimulus level 7) boundaries, averaged across both CW and CCW orientations. Lower values indicate a bias to perceive stimulus orientations closer to the more horizontal boundary. Right panel: Mean RT (\pm SEM) expectation effect during the first part of the test session. * $p < 0.05$. (B) Expectation effects (Unexpected – Expected) plotted with raincloud plots (Allen et al., 2019) displaying probability density estimates (upper) and box and scatter plots (lower), for PSEs (top) and RTs (bottom). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive PSE expectation effect values indicate perceived orientations closer to the more horizontal boundary for expected events relative to unexpected events, while positive RT effects reflect faster responses on expected trials..... 101

Figure 4.4. Experimental paradigm. (A) A centrally presented visual cue instructed participants to abduct either their index or little finger. Each finger abduction predicted an oriented stimulus and participants were required to respond (yes/no) to whether the stimulus was clockwise or counter-clockwise. In training, actions perfectly predicted the stimulus orientation that was presented. (B) In the scanning session 24 hrs later, participants completed the same task but actions now predicted the oriented stimulus with 33.3% validity, presenting expected and unexpected stimuli 33.3% of the time each and on the remaining trials the stimulus was subsequently omitted..... 110

Figure 4.5. (A) Sagittal view of an example mean functional image of the calcarine sulcus and occipital lobe overlaid with the estimated V1 ROI (in white). (B) A schematic representing classification of CW and CCW activity patterns across voxels using MVPA, adapted from Haynes and Rees (2006). (C) Mean decoding accuracy (\pm SEM) for expected and unexpected events, showing superior classification of stimulus orientation when stimuli were expected. (D) Mean BOLD activity (\pm SEM) for voxels

tuned towards and away from the presented stimulus separately for expected and unexpected events. Expectation suppression is only observed in voxels tuned away from the presented stimulus (** $p < .01$, *** $p < .001$). (E) Expectation effects (Expected BOLD – Unexpected BOLD signals) were calculated for voxels tuned towards [teal], and away from [blue] the presented stimulus and are plotted with raincloud plots (Allen et al., 2019) displaying probability density estimates (upper) and box and scatter plots (lower). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=20). Negative effects indicate lower BOLD signal for expected events relative to unexpected events..... 115

Figure 5.1. Experimental paradigm. (A) A centrally presented visual cue instructed participants to abduct either their index or little finger. In training, each finger lift perfectly predicted an oriented stimulus and participants were required to respond (yes/no) to the question in the middle of the screen – indicating whether the stimulus was clockwise (‘CW?’) or counter-clockwise (‘CCW?’). (B) In the test session, participants performed the same task but were allocated into one of three groups where actions now predicted stimulus orientations with either 50%, 66.7% or 75% validity..... 129

Figure 5.2. (A) Mean RTs (\pm SEM) were faster for expected than unexpected trials in both 66% and 75% contingent groups but not when 50% contingent. (B) Mean proportion of errors (\pm SEM) were lower for expected than unexpected trials in both 66% and 75% contingent groups but not when 50% contingent (** $p < .01$, *** $p < .001$). (C) RT expectation effects (Unexpected RT – Expected RT) were calculated for the 50% [peach], 66% [teal], and 75% [blue] contingent data and are plotted with raincloud plots (Allen et al., 2019) displaying probability density estimates (upper) and box and scatter plots (lower). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N = 24 per group). Positive effects indicate faster responses to expected events relative to unexpected events. (D) Mean RT effects (Unexpected RT – Expected RT, \pm SEM) by mini-block across the test session..... 132

Figure 5.3. (A) Schematic representation of proposed extrinsic feedforward (red) and feedback (blue) connections across layers between the LGN, V1 and V2; taken from Lawrence, Formisano, et al. (2019). (B) Hypothesised orientation-specific BOLD

(preferred – non-preferred) responses to expected and unexpected events across layers, and according to key assumptions of predictive coding theories. Arrows highlight a possible reconciliation to the paradox that expectation suppression and enhancement (or ‘sharpening’) are represented in different cortical layers..... 141

Figure 5.4. Experimental paradigm. (A) The same trial sequence as Experiment 6. A centrally presented visual cue instructed participants to abduct either their index or little finger. Each finger abduction predicted an oriented stimulus and participants were required to respond (yes/no) to whether the stimulus was clockwise or counter-clockwise. In training, actions perfectly predicted the stimulus orientation that was presented. (B) In the scanning session 24 hrs later, participants completed the same task but now actions predicted the oriented stimulus with 75% validity..... 143

Figure 5.5. (A) Visualisation of the selected anatomical V1 ROI (light grey) on a mean functional image of an example participant. Overlaid red and yellow lines represent co-registered anatomical WM (yellow) and pial surface (red) boundaries to the mean functional image. The left panel shows voxels that were significantly active to the presented stimuli in the functional localiser task (green). The right panel shows the voxels within these active voxels that were strongly tuned to CW (pink) or CCW (blue) orientations. (B) A mean functional image overlaid with distributions of voxels in superficial (green), middle (blue) and deep (red) layers of the GM. (C) A schematic representing the spatial regression approach used to determine the volume distribution of a selected voxel (e.g., red square) over superficial (I), middle (II), and deep (III) cortical layers to estimate laminar time-courses. Taken from Aitken et al. (2020) 149

Figure 5.6. (A) Orientation-specific BOLD responses in V1 across cortical layers (deep, middle, and superficial) for expected (blue) and unexpected (orange) events. The same effects are shown averaged over all participants (left panel) as well as presenting individual participants’ data points (right panel). (B) Presents the same as (A) but for orientation-specific BOLD responses in V2. Error bars indicate within-subject SEM 151

Figure 5.7. (A) BOLD responses in V1 across cortical layers for expected (left panel) and unexpected (right panel) events. Within each layer, responses are separated according to voxels preferring (solid lines) or non-preferring (dashed lines) the

presented stimulus orientation, and demonstrates global BOLD signal before orientation-specific subtractions. Violin plots show individual data points for preferring (filled in) and non-preferring (outlined) voxels across layers (deep, middle and superficial). Error bars indicate within-subject SEM..... 154

Chapter 1: Introduction

The brain is constantly bombarded with an array of noisy sensory information that needs to be combined and interpreted in order to construct representations of the world. Integration of these sensory signals is often imperfect due to signal transmission in the brain but also uncertainty in the information reaching our receptors (Press et al., 2020b). Consider for example, that you are driving in heavy rainstorm; the visual and auditory input of large and rapid raindrops hitting your windscreen is likely to mask your view of the road as well as the voice of your friend in the passenger seat giving you directions. Perception is therefore not an easy task for the brain, as it needs to be able to make sense of ambiguous input rapidly if it is to be of use when interacting with our environment e.g., being able to perceive the left turning in the road and act accordingly before driving past.

Our sensory system therefore needs to overcome two main challenges in the face of such noise; firstly representations need to be constructed that are veridical – therefore reflecting accurately what is really there – and secondly informative – to inform belief updating when we are presented with surprising information (Press et al., 2020b). One way that these perceptual challenges may be overcome is by using our prior knowledge about what we are likely to perceive, based on what we have perceived in the past, to make inferences about noisy sensory input. This idea originated through Helmholtz’s notion of ‘unconscious inference’ (von Helmholtz, 1867) and has had a highly influential impact across cognitive psychology in the last century.

The notion that we can use prior knowledge to infer properties of objects and scenes in ambiguous input is thought to explain a variety of popular visual illusions. For example, we often perceive concave faces as convex because our experience of viewing faces in the natural world is typically with this convex structure (see Fig. 1.1.). These types of

illusions likely occur because ambiguous shadow information (Gregory, 1997) is combined with knowledge learned over the lifespan of how lighting typically frames faces as they appear convex (de Lange et al., 2018). The brain then infers the most likely interpretation of the input based on this prior experience with the world. Additionally, priors about the typical dimensionality of objects in the world can help us to infer 3D representations of shapes from simple 2D wire-frame visual inputs e.g., perceiving a 3D rhomboid shape in the Necker cube illusion (Gregory, 1997). Finally, perception can be improved by prior knowledge by ‘explaining away’ expected features of a scene in order to highlight those unexpected features that might be vitally important to perceive e.g., an unexpected tiger hidden amongst contextually expected stripy leaves (Kok & De Lange, 2015, see Fig. 1.1.).

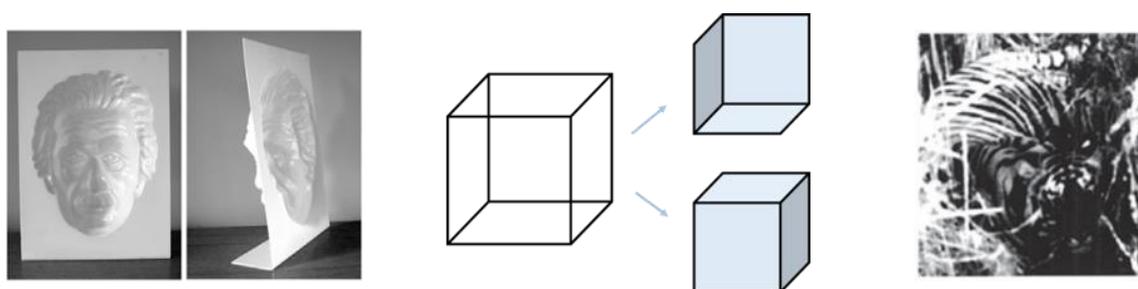


Figure 1.1. Left: Hollow-face illusion, concave faces are perceived as convex when full shading cues are available (e.g., front on; left), adapted from Press et al. (2020b). Middle: the Necker cube can be perceived in two different 3D perspectives. Right: after ‘explaining away’ the leaves in ambiguous input the tiger stands out more (taken from Kok & De Lange, 2015).

1.1. What are ‘top-down’ effects and how is ‘prediction’ defined?

One way that psychologists often describe the process of using prior knowledge to help shape perception and overcome computational challenges, is to describe such influences as ‘top-down’ processes. This is in contrast to ‘bottom-up’ processes that do not rely on any internal knowledge but instead rely purely on external input from our environment

(Gibson, 1972). The metaphors “top” and “bottom” are thought to reflect an assumed hierarchy of information processing, such that information that starts as simple retinal input (“bottom”) is passed up to successively higher-order cognitive representations (“top”; see Shea, 2014). The distinction has also however been established by neuroscientists to reflect message passing throughout the cortex according to a hierarchical organisation of functions (Friston, 2008; Mechelli et al., 2003). For instance, recent popular theories propose that ‘top-down’ processes are analogous to descending feedback connections from ‘higher’ level areas (e.g., frontal) to deep cortical layers of ‘lower’ level areas (e.g., early visual), while ‘bottom-up’ processes convey ascending feedforward activity from the earliest visual areas to middle and superficial layers of ‘higher’ level areas (Bastos et al., 2012).

Regardless of the precise distinction, ‘top-down’ effects on perception have formed a foundation in sensory cognition literatures for research on how predictions (or ‘expectations’) influence our perceptual experience of the world. Prior knowledge about the statistical regularities of our environment are learned over our life-span e.g., knowledge that light often comes from above helps us to predict object shapes on the basis of shading patterns (de Lange et al., 2018), and is often taken to define how we might form ‘predictions’ about incoming sensory input. Knowledge of the statistical regularities of a given context also help us to form predictions about what we should ‘expect’ to be there in a certain scene based on the surrounding properties, and this is thought to aid object recognition (Bar, 2004).

Predictions that we form based on such experiences are likely to be an accurate reflection of the state of the world considering that our sensory environment remains relatively stable over time, meaning that we can build internal models based on past experiences that should also accurately reflect present and future incoming sensory

input too. A key aspect of this definition of prediction, is that predictions should be formed on the basis of contingency and contiguity between sensory states or actions (Press et al., 2020b), e.g., pressing a doorbell button typically elicits the auditory input of a bell ringing and visual input of a door opening and seeing our neighbour. These kinds of predictions are formed through statistical learning – a process that is thought to be often implicit and by which we are able to extract statistical regularities of co-occurring events in our environment (Schapiro & Turk-Browne, 2015; Turk-Browne et al., 2010), and occurs continuously across our lifetime (Campbell et al., 2012; Kirkham et al., 2002). This definition of prediction – according to learned statistical probabilities of our environment – is the one adopted in this thesis.

It is also important to note that there are a range of other top-down influences on perception which do not necessarily rely on prediction. For example, it has been suggested that top-down attention and top-down prediction can be differentiated according to task relevance or statistical likelihood of occurrence, respectively (Summerfield & Egner, 2009). However, in our natural environment as well as the majority of experimental paradigms, more probable events are also more relevant for task performance. For example, many classic attentional paradigms in fact manipulate stimulus probabilities (Posner et al., 1980), while expectation manipulations render the relatively more expected also more task relevant (Rohenkohl et al., 2014). It is therefore important to orthogonalise predictive and attentional influences with definitions of stimulus probability and task relevance, respectively.

The focus of this thesis is to examine how sensorimotor predictions influence perception, and therefore the definition of prediction, as outlined according to stimulus probabilities, is adopted throughout this thesis in order to isolate these kinds of predictive influences on perception.

1.2. How do predictions influence perception?

So far, some different types of top-down influences have been discussed, introducing the idea that predictions may have a profound influence on perception. This section will now discuss *how* or in what ways predictions influence perceptual processing.

1.2.1. Predictions enhance perception

One way in which our perceptual systems are thought to allow for generation of appropriate experiences from the vast array of ambiguous input and sensory noise is through Bayesian inference. Bayesian accounts of perception have origins in Helmholtz's theory of unconscious inference through the idea that our sensory systems use prior knowledge to resolve sensory ambiguity, by inferring the most likely properties of our environment (Kersten et al., 2004). Such accounts are popular in the normative sensory cognition literature, and suggest that we combine our expectations (prior) with the input (likelihood) to determine what we perceive (posterior; de Lange et al., 2018; Kersten et al., 2004; Knill et al., 1996). In other words, these theories propose that we integrate our prior knowledge into our representations of the world, using knowledge about the common statistical regularities in our sensory input to bias our perceptual experiences *towards* those that are more probable (Kersten et al., 2004; Kersten & Yuille, 2003; Knill et al., 1996). That is, we *upweight* perception of expected events, increasing the detectability and apparent intensity of such events (Brown et al., 2013; Carrasco et al., 2004; Wyart et al., 2012) and enabling rapid generation of largely veridical experiences in the face of sensory noise (de Lange et al., 2018; Kersten et al., 2004; Yuille & Kersten, 2006). For example, perceiving what we expect would help us to understand what a friend is saying in a loud and crowded bar. These mechanisms are thought to be adaptive because they will generate experiences that, on average, are more veridical and will more accurately depict our sensory environment.

Such accounts are consistent with a range of findings that expected visual events are recognised and identified more readily (Bar, 2004; Palmer, 1975; Puri & Wojciulik, 2008), and perceived to be more intense (Han & van Rullen, 2016), than unexpected events. For example, we are faster to identify a loaf of bread in the expected context of a kitchen counter, than an unexpected mailbox on the same counter (Palmer, 1975). Other findings supporting these accounts have shown that integration of expected (congruent) sensory information can facilitate detection and identification of objects (Noppeney et al., 2008), as well as enhancing motion discrimination (Lewis & Noppeney, 2010). Studies examining perceptual sensitivity to visual features in signal detection paradigms are also in-line with upweighting theories, which have shown that participants are biased towards perceiving expected stimuli (Wyart et al., 2012), and have heightened sensitivity for expected orientations, and even for features that represent more extreme exemplars of the expected category (Cheadle et al., 2015).

Furthermore, studies have shown that such effects likely result from learned models of the world built upon statistical regularities. For example, in a dot motion task where certain motion directions are presented a higher proportion of the time, perception is found to be biased towards these highly probable motion directions and this bias is especially strong on signal-absent trials (Chalk et al., 2010). Additionally, stimuli that are cued with high validity have been found to enter conscious awareness faster than stimuli that are invalidly cued during a continuous flash suppression task (Pinto et al., 2015). Finally, predictions about upcoming stimuli have been demonstrated to decrease response times and increase accuracy for expected events relative to unexpected or unpredictable events, which is thought to reflect highlighted processing of the expected (e.g., Kok et al., 2014; Kok, Rahnev, et al., 2012; Richter & de Lange, 2019). In sum, these findings in the broader sensory cognition literature are consistent with

assumptions of predictive upweighting, whereby expectations are integrated with sensory input to enhance perception and aid veridicality.

1.2.2. Predicted action outcomes are perceptually attenuated

The evidence discussed up until now has demonstrated that sensory events can be predicted by other sensory events, yet such sensory outcomes can also be predicted by action. While upweighting theories of prediction provide an appealing explanation for how we generate veridical perceptual experiences in a noisy sensory world, they contrast with prominent theories of predictive influences in action domains.

In action domains, prominent cancellation models propose that a forward model of the motor system predicts the sensory consequences of our actions (Wolpert & Flanagan, 2001). Under these accounts, such predictions – or ‘efference copies’ thought to simulate the sensory consequences of action – are purported to be subtracted from subsequent sensory input (Wolpert & Flanagan, 2001). This process results in attenuated – or *downweighted* – perception of events matching those predictions (Blakemore et al., 1998; Dogge, Custers, & Aarts, 2019; Fiehler et al., 2019; Kilteni & Ehrsson, 2017a). In a limited capacity system, such mechanisms are assumed to prioritise processing of unexpected sensory events which are more informative to the individual – signalling the need for corrective action or updating of our models of the world (Press et al., 2020b; Wolpert & Flanagan, 2001). For example, before picking up a cup of tea we can predict certain properties of the cup, such as its weight – using our available sensory input e.g., the thickness of rim and how full it is – and use these predictions to adopt the appropriate action and force to pick it up. However, if the cup of tea is lighter than expected, attenuated processing of expected signals (e.g., touch on our fingertips) will allow dedicated processing of unexpected events (e.g., accelerating motion of the cup), allowing swift updating of our beliefs about the environment (e.g.,

the weight of the cup) and support corrective action to avoid dropping or spillage. These mechanisms are thought to be adaptive because they will generate experiences that are more informative to the organism and therefore aid in rapid updating of our models of the ever-changing world.

Support for this account is vast in humans as well as in a variety of other species and sensory systems. For example, Mormyrid fish can improve their detection of prey-like stimuli by attenuating input of internally-generated electric fields – a finding thought to be a consequence of being able to predict self-generated sensory input (Enikolopov et al., 2018). Similarly, mice trained in virtual reality show suppressed auditory responses to self-produced tones generated by treadmill running (Schneider et al., 2018) or licking behaviours (Singla et al., 2017), compared to no movement. In humans, much of the evidence supporting this account is demonstrated in the force judgement paradigm whereby participants perform actions with their right hand to produce a force to their stationary left hand. These studies report that self-produced tactile sensations are perceived to be less intense than the same forces that are ‘externally-produced’, for example by a robot or computer (Bays et al., 2005, 2006; Blakemore et al., 1998, 1999; Kiltani et al., 2019; Kiltani & Ehrsson, 2017a; Wolpe et al., 2016, 2018). These findings are often interpreted to result from a predictive attenuation mechanism on the premise that we are able to predict the consequences of our own actions but not similar externally generated events. Such studies and conclusions have also provided an explanation for why it is difficult to tickle oneself (Blakemore et al., 2000), since attenuating the consequences of our actions would reduce the perceived intensity of sensations caused by intended self-tickle.

Recent studies have demonstrated that tactile attenuation can be observed in conditions where participants do not perform actions, but instead when a rubber hand, believed to

be one's own hand, performs the action producing tactile events (Kilteni & Ehrsson, 2017a). Additionally, tactile attenuation can be observed when participants' right-hand action applies touch to the left hand with a tool (Kilteni & Ehrsson, 2017b). Although forces were not generated by participants' own effectors, it suggested that these findings support cancellation mechanisms because the brain forms internal representations of body ownership (e.g., rubber hand in the former) and of the current effector (e.g., extension to the tool in the latter), which provides input to the forward model generating predictions about the resulting force. Other recent studies have provided support for cancellation models by demonstrating that tactile attenuation is strongest for reaches towards other body parts relative to reaches towards non-body-related external targets (Gertz et al., 2017).

Although largely in the tactile domain, support for these theories is also noted in visual and auditory sensory modalities. For example, studies show reduced perceptual sensitivity to visual action outcomes relative to stimuli not associated with action (Cardoso-Leite, Mamassian, Schuttz-Bosbach, et al., 2010) and attenuated perceived loudness of self-generated sounds relative to externally-generated sounds (Weiss et al., 2011). Additionally, studies that have trained participants that left and right hand actions predicted particular letter stimuli (Roussel et al., 2014) or visual Gabor patches (Desantis et al., 2014), revealed worse contrast discrimination and reduced sensitivity to tones accompanying visual events for expected relative to unexpected action-outcomes, respectively. Assumptions that action predictions have a distinct attenuating influence on perception has further advanced theoretical accounts of clinical disorders and diseases including those of hallucinations in schizophrenia, Parkinson's and Alzheimer's (Corlett et al., 2019). Furthermore, studies adopting the typical tactile action set-ups described, have demonstrated sensory differences associated with healthy

ageing (Wolpe et al., 2016), motor severity in Parkinson's disease (Wolpe et al., 2018) and reduced tactile attenuation in patients with schizophrenia (Shergill et al., 2005).

Taken together, these findings in action domains typically suggest that expected action outcomes are perceptually attenuated relative to outcomes that are externally-generated, and are assumed to operate as a function of cancelling predicted consequences of action from perception.

1.2.3. Do action predictions influence perception in a distinct fashion?

The downweighting influence of prediction in action domains is therefore in apparent contrast to upweighting influences in broader sensory cognition domains (note that attenuation and downweighting, as well as enhancement and upweighting, will be used interchangeably throughout this thesis). The primary question of this thesis is thereby to investigate these paradoxical influences, and more specifically, examine whether action predictions indeed influence perception in a distinct fashion to other types of predictions outside of action, when comparable methodologies and manipulations are adopted.

Although both processes appear to have adaptive advantages, it is uncertain why the adaptive arguments presented for downweighting (informativeness) and upweighting (veridicality) predicted percepts should apply differentially across domains, because it appears just as crucial to optimise informativeness and veridicality regardless of how we form predictions (Press et al., 2020b). To begin with, I will consider whether there is in fact no paradox at all because predictive mechanisms need to subserve different functions in the context of action. Downweighting mechanisms in action are thought to favour perceptual informativeness because it is essential to be able to guide corrective actions and update our models of the world. For instance, we would likely be slower and less effective in responding to unexpected properties of objects we interact with if priority was given to processing of expected features instead, e.g., there would be a

higher likelihood of dropping a surprisingly light cup of tea. Additionally, downweighting mechanisms are perhaps less essential in non-active perception as there is no need for rapid corrective action. These opposing theoretical accounts may therefore be reconciled by a conclusion that sensorimotor predictions influence perception differently from other types of prediction. Furthermore, if prediction mechanisms subserve a distinct function during action, evidence for downweighted perception of predicted action outcomes should predominate across action literatures.

However, taking a closer look at the supporting literature across the two fields, it is notable that there are very different types of evidence supporting each theory, which makes comparisons across the two fields difficult. In particular, the majority of evidence supporting cancellation theories comes from action paradigms examining the perceived force of suprathreshold tactile events to the finger tips (e.g., Bays et al., 2005; Blakemore et al., 1998; Kiltner & Ehrsson, 2017b), whereas a large amount of literature supporting upweighting theories in sensory cognition domains often examines sensitivity towards detecting ambiguous visual events (e.g., Cheadle et al., 2015; Wyart et al., 2012). This distinction is potentially imperative to understanding the paradox considering that stimuli that are presented at-threshold are ambiguous in nature and are therefore often barely detectable, perhaps highlighting the need to upweight perception towards what we expect to form accurate perceptual representations. On the other hand, perception of suprathreshold stimuli that are highly detectable may reveal downweighting effects that only operate when highly surprising unexpected events are presented (Press et al., 2020b). There is some preliminary evidence in favour of this proposition which demonstrates that the perceived loudness of self-generated auditory events is attenuated relative to passive events when presented with suprathreshold

intensity, yet self-generated events presented at-threshold were perceived louder than equivalent passive events (Paraskevoudi & SanMiguel, 2020; Reznik et al., 2015).

These findings provide preliminary support for the possibility that methodological differences across the two domains contribute to contrasting findings. However, an additional factor that makes comparisons difficult relates to the fact that the evidence tends to focus on different sensory modalities – namely, touch in action domains and vision outside of action. Despite the sparsity of converging evidence across domains, some recent studies have examined predictive influences on *visual* action outcomes, providing some insight to the comparability across sensory domains. These studies revealed that participants tend to visually perceive congruent finger movements *more* intensely (Yon & Press, 2017, 2018) and are biased towards reporting their presence (Yon et al., 2020), than incongruent outcomes. Additionally, participants are reported to be better at detecting biological motion congruent with action (Christensen et al., 2011), more often perceive an ambiguous rotating sphere rotating in the same direction as their action (Dogge, Custers, Gayet, et al., 2019), and have enhanced visual discrimination performance during related grasping (Guo & Song, 2019). Furthermore, participants have been reported to have biased visual estimates of hand location in line with the expected trajectory of action (McDonough et al., 2019), that is greater with more predictive (e.g., 75% validity) action intentions, and therefore trajectories (Hudson et al., 2017). These findings suggest that action predictions perhaps do not always exert an attenuating influence on perception and demonstrate that visual perception can be biased *towards* expected action outcomes – more consistent with upweighting, than downweighting theories of perception. Such findings however contrast with some evidence from a popular study demonstrating attenuated perceptual sensitivity to visual action outcomes (Cardoso-Leite, Mamassian, Schütz-Bosbach, et al., 2010), but it is

worth noting that a later study reported being unable to replicate this effect and suggested that it may instead reflect mechanisms that are not predictive (Schwarz et al., 2017).

Taken together, preliminary evidence from these studies suggest that downweighting mechanisms may not subserve the same predictive function in visual perception as has been demonstrated across a wider range of touch studies. Therefore, due to the large disparity between methodological approaches and sensory modalities, a conclusion that action predictions have a distinct functional influence on perception would be premature on the basis of these findings.

1.2.4. Do action predictions influence tactile perception in a special way?

As suggested, a stark difference between studies purporting to demonstrate upweighting and downweighting is that the former study visual perception whereas the latter study tactile perception, even within action domains. A possible alternative resolution to the opposing findings and theoretical assumptions is that predictive downweighting mechanisms – thought to be specific to action – may in fact be specific to tactile perception, and indeed there is a much larger consensus that predicted tactile action-outcomes are attenuated.

Accordingly, many researchers believe that action predictions shape tactile perception in a qualitatively distinct way, with proposals that differences from other senses relate to tactile events being body-related (Dogge, Custers, & Aarts, 2019) in a way that many predicted visual or auditory events are not. For instance, somatosensation is thought to be unique in the way it acquires its inputs, through tight coupling with the motor system, relative to vision and audition (Kusnir et al., 2019). As such, it could be argued that rapid detection of an unexpected tactile event on the skin, e.g., a crawling and

potentially poisonous spider, is likely to be more important for immediate action than unexpected visual events e.g., that same spider hanging on its web in front of you.

However, if action indeed requires distinct prediction mechanisms that fulfil the need to update models of the world and aid corrective actions, it would perhaps seem strange that these adaptive mechanisms should work differently across sensory modalities. Such claims would be inconsistent with proposals that predictive mechanisms should operate similarly across sensory modalities that contain attributes of stimulus intensity (Brown et al., 2013), and begs the question of why other sensory consequences should not require the same adaptive processes within action. There is limited evidence for answering such a question however, and the lack of comparability across sensory modalities and domains means that the complexities of this debate remain unsolved without further examination. It is quite possible that differences in methodologies adopted across the two domains or sensory modalities contribute to the apparent paradoxical influences of prediction on perception in the literature, and this thesis therefore aims to determine these precise influences when manipulations are comparable to broader sensory cognition domains.

1.2.5. Open questions

As outlined, there is conflicting evidence for both upweighting and downweighting influences of prediction on perception across action and broader sensory cognition domains that notably adopt quite diverse methodologies. It is therefore unclear whether there are distinct adaptive mechanisms operating in action contexts that produce opposing perceptual effects to influences of prediction outside of action, or whether the differences across methodologies muddy the water on a resolution of this debate. The first part of this thesis therefore aims to examine this question directly by investigating the influences of sensorimotor prediction on touch perception. First, it aims to replicate

a typical tactile attenuation finding using a force judgement paradigm, while adding an additional condition to control for a possible methodological confound in present paradigms (Chapter 2). It subsequently combines this set-up with a paradigm that manipulates prediction according to statistical regularities, as is typical in sensory cognition domains (Chapter 3). These experiments aim to determine whether action predictions exert a distinct downweighting influence on touch perception, or whether the mechanisms operating can generate upweighting influences comparable to those in the sensory cognition literature, when adopting similar manipulations of predictability.

1.3. The predictive nature of cortical effects – does the perceptual paradox extend to neural mechanisms?

The behavioural findings discussed so far indicate a paradox between influences of prediction on perception in action and broader sensory cognition domains. For the second part of this thesis, I will consider whether this paradox extends to neural prediction mechanisms in sensory cortices and whether there is any further clarity to the resolution of this debate by examining how neural findings fit with the two theories.

1.3.1. Suppressed neural response to predicted consequences of action

The cancellation account is also supported by a variety of neural studies that find reduced activity in sensory brain regions for predictable action-outcomes relative to externally generated outcomes. For example in touch, studies investigating predictive mechanisms have found reduced neural response in the secondary somatosensory cortex and cerebellum to tactile forces produced by one's own action relative to externally generated forces – in similar force judgement paradigms as described earlier (Blakemore et al., 2001; Kilteni & Ehrsson, 2020a; Shergill et al., 2013; Wolpe et al., 2018). Similarly, in vision, a study by Stanley and Miall (2007) revealed reduced Blood Oxygenation Level Dependent (BOLD) signals in visual cortices to observed visual

actions that were congruent with participants' executed actions relative to incongruent actions. Comparable effects have also been observed in speech perception, where the production of inner speech that matched a simultaneous audible phoneme was associated with suppressed electrophysiological activity, relative to a mismatching phoneme (Whitford et al., 2017).

Such effects are sometimes more broadly referred to as expectation suppression, defined as a reduction in a measure of neural activity associated with an expected sensory event compared to events that are unexpected. Studies demonstrating suppression effects in action domains are therefore often proposed to provide supporting evidence for the operation of predictive downweighting mechanisms in sensory cortices. Under the cancellation account, a reduction in neural response is proposed to result from predictions – about the sensory consequences of action – being sent to the respective sensory brain regions and then subtracted from the sensory input, resulting in suppression of expected sensory units, and allowing us to become sensitive to unexpected events (Wolpert & Flanagan, 2001).

It is unclear whether or how findings of reduced neural response may be associated with findings of reduced perception of expected action outcomes, especially considering the differing types of evidence. With behavioural attenuation studies in action forming the majority of evidence in favour of cancellation accounts, while evidence outside of action presents opposite upweighting effects, interestingly, this pattern is not equivalent amongst neural evidence. Perhaps surprisingly, expected events outside of action contexts are not often associated with an increase in BOLD activity, but instead with a reduction in neural signal (e.g., Alink et al., 2010; Han et al., 2019; Kok, Jehee, et al., 2012; Richter et al., 2018; Todorovic & De Lange, 2012). While these findings may suggest some similarities with how prediction mechanisms are neurally implemented

across domains, the precise mechanisms underlying expectation remains controversial and will be discussed in the following section.

1.3.2. Sharpening of expected sensory representations

Although there is evidence for neural expectation suppression in both action and broader sensory cognition domains, some studies outside of action have argued that the lower overall signal for expected events is not necessarily reflective of a predictive downweighting mechanism – as action theorists would suggest. For example, Kok and colleagues (2012) trained participants that different pitched tones could predict grating orientations with 75% validity in order to examine the activation patterns in Primary Visual Cortex (V1) to expected or unexpected stimuli. The authors found that expected orientations were associated with a lower BOLD signal than unexpected orientations, however, they also trained pattern classifiers to discriminate between stimulus orientations and found that these classifiers performed with greater accuracy when decoding expected orientations (Kok, Jehee, et al., 2012). In other words, the multivariate pattern analysis (MVPA) suggested that there was improved representation of expected events, such that more information about stimulus orientation could be read out from the signal in V1 when the event was expected.

The authors suggested that these results could be explained by an underlying ‘sharpening’ mechanism – in which it is proposed that predictions increase the gain on sensory populations tuned towards features or stimuli that are expected, and via lateral inhibition relatively suppress processing in sensory populations tuned to unpredicted or irrelevant features (de Lange et al., 2018; Kok, Jehee, et al., 2012; Press et al., 2020b; Summerfield & de Lange, 2014; Yon et al., 2018). Put another way, expectation suppression will only be seen in neuronal populations tuned away from the expected stimulus (e.g., populations tuned to vertical orientations when horizontal orientations

are expected), while populations tuned towards expected events retain their activation or are even enhanced – resulting in a ‘sharpened’ representation of the expected stimulus (de Lange et al., 2018; Han et al., 2019; Kok, Jehee, et al., 2012; see Fig. 1.2.). This pattern may plausibly underlie other expectation suppression findings across domains, since only a very finite population with the highest selectivity may be enhanced while the remaining populations tuned to unpredicted stimuli will be suppressed, thus resulting in a larger, more detectable, general suppression effect.

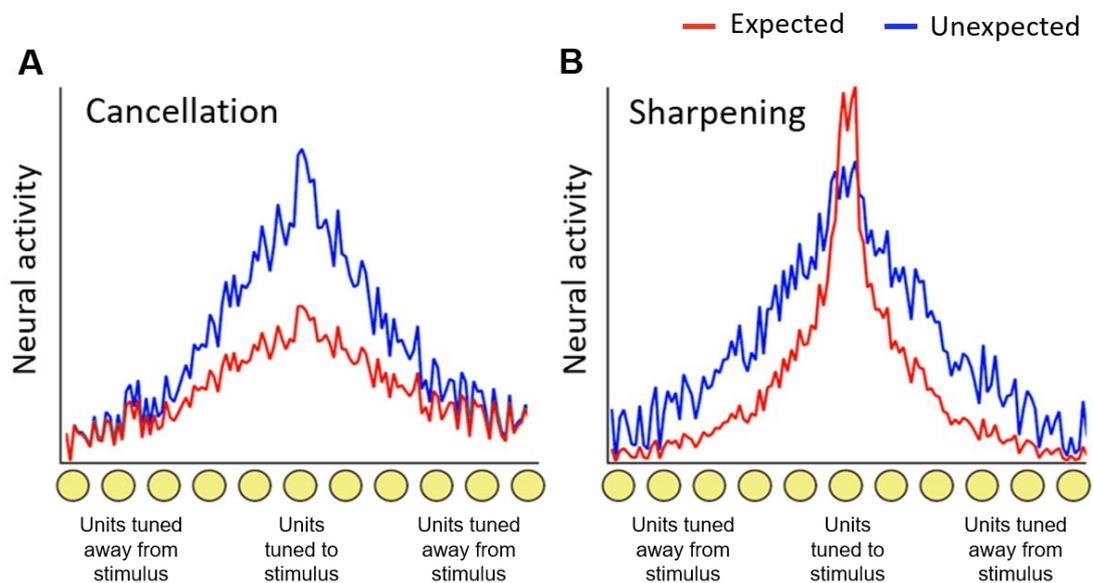


Figure 1.2. (A) Cancellation theories propose that neural activity associated with expected outcomes of action is suppressed in order to highlight unexpected events, resulting in expectation suppression effects that are greatest in populations tuned towards the presented stimuli. (B) Sharpening theories propose that we ‘pre-activate’ populations tuned to the expected stimulus and relatively suppress processing in populations representing other events, resulting in a sharpened neural response (adapted from Press et al., 2020b).

Such sharpening accounts are more in line with perceptual upweighting theories which propose that expectations formed from our knowledge about the statistical regularities of the world bias our sensory representations, altering the weight on different sensory

channels and effectively ‘turning up the volume’ on expected inputs (de Lange et al., 2018; Press et al., 2020b; Yon et al., 2018). A number of neuroimaging techniques have additionally provided evidence for such a mechanism operating in broader sensory cognition domains. For instance, transcranial magnetic stimulation demonstrates that disruptive pulses to sensory regions at the time of predictive cues eliminates perceptual cueing effects (Gandolfo & Downing, 2019), magneto- and electro-encephalography (M/EEG) work demonstrates that expected visual events can be decoded from visual brain activity ahead of their expected presentation (Blom et al., 2020; Kok et al., 2017) and fMRI studies indicate that expectations can induce stimulus-specific templates with no bottom-up input (Aitken et al., 2020; Ekman et al., 2017; Kok et al., 2014).

These studies are thought to provide support for a possible mechanism of upweighting accounts that predictions can ‘pre-activate’ sensory representations such that evidence for expected representations can be observed even before input is expected, or in the absence of bottom-up input. Specifically, expected stimulus orientations can be decoded 40 ms prior to sensory input and provide evidence for the idea that expected stimulus templates are pre-activated prior to expected input (Kok et al., 2017). Similarly high-resolution EEG methods have been used to show that signals of moving objects can be decoded ahead of predictable stimulus trajectories, and importantly in positions where a stimulus was never present (Blom et al., 2020). Additionally, when macaques are presented with pairs of stimuli in an associative learning task, predictable images have been shown to elicit higher firing rates of temporal cortex neurons ahead of presentation, and especially those tuned to expected stimuli (Erickson & Desimone, 1999). Finally, studies that measure activity patterns on trials when events are predicted by imperative cues but do not subsequently occur can provide evidence for activation of expected representations – finding similar patterns of activity for expected stimuli in

visual areas on omission trials to patterns evoked by presented expected stimuli (Kok et al., 2014). Other paradigms have demonstrated evidence for stimulus-specific templates in non-stimulated visual regions using set-ups that occlude portions of naturalistic scenes, and using MVPA show that information in these missing quadrants could be decoded from V1 (Smith & Muckli, 2010).

Interestingly, these studies and the majority of neural work supporting the operation of neural sharpening mechanisms comes from outside of action contexts, which again raises the possibility that expected sensory representations may be influenced in a distinct fashion in action contexts relative to broader sensory cognition domains. Consistently, propositions that action predictions – especially in touch – influence sensory representations distinctly from other types of prediction may relate to assumptions that tactile attenuation during action is dependent upon somatosensory-cerebellar connectivity (Blakemore et al., 1998; Kilteni & Ehrsson, 2020a), in contrast with hippocampal mediation of prediction in visual processing (Kok et al., 2020; Kok & Turk-Browne, 2018). However, it is also possible that findings of expectation suppression observed in both domains are instead explained by underlying sharpening mechanisms, which would suggest a more domain general influence of prediction on perceptual processing. For example, it is possible that findings of reduced somatosensory (Blakemore et al., 1998; Kilteni & Ehrsson, 2020a; Shergill et al., 2013) or visual (Kontaris et al., 2009; Stanley & Miall, 2007) neural response when perceiving expected action outcomes – frequently taken to support downweighting theories – may plausibly reflect the operation of predictive upweighting or ‘sharpening’ mechanisms when the appropriate analytical methods are adopted to detect such influences.

1.3.3. Open questions

While cancellation or ‘downweighting’ theories can be supported by evidence of expectation suppression in both action and broader sensory cognition domains, some initial evidence suggests that these univariate suppression effects may instead be reflective of underlying neural sharpening mechanisms. This is an interesting possibility, considering that many researchers believe that action predictions influence perception and sensory processing in a distinct way from other types of predictions, and much of the neural support for cancellation theories comes from findings of reduced BOLD signal for predicted consequences of action. The second part of this thesis therefore examines the underlying neural mechanisms of predictions formed on the basis of action, in order to determine whether they operate similarly to domains outside of action or whether there may be evidence for distinct underlying mechanisms. This question is investigated by manipulating the statistical relationship between action and visual outcomes – similarly to in Chapter 3 and literature examining predictions outside of action contexts – and uses both univariate and multivariate analysis techniques to explore the possibility that sharpening mechanisms are underlying effects of expectation suppression in action (Chapter 4).

1.4. Could distinctly operating prediction mechanisms across cortical layers reconcile the neural paradox?

1.4.1. Hierarchical theories of prediction in the cortical column

The human brain consists of cerebral cortex (or grey matter) that serves as the outer most layer of neural tissue of the cerebrum. This layer is approximately 3 mm thick and can be further broken down into six cytoarchitectonic layers (Brodmann, 2007), that can be distinguished by their different cell types and are thought to have different functions. For instance, the cortex is believed to be hierarchically organised, in which extrinsic

feedforward connections relay information originating in layers 2/3 (superficial) of lower cortical (e.g., V1) or sub-cortical (e.g., Lateral Geniculate Nucleus (LGN) or thalamus) regions to layer 4 (middle layer) populations within higher areas (e.g., frontal), while feedback connections relay information from deep layers 5/6 of higher level cortical regions to the agranular (superficial and deep) layers of lower level regions (Bastos et al., 2012; Felleman & Van Essen, 1991). The precise functional roles of these ascending and descending connections across the hierarchy are however topics of debate.

One theory, namely predictive coding, suggests that the role of feedback connections is to mediate contextual guidance to lower level inputs by forming predictions about our incoming sensory input, while the role of feedforward connections relates to the segregation of sensory information (Friston, 2005). More specifically, predictive coding accounts suggest that in each level of the cortical hierarchy there are separate sub-populations of neuronal units that encode not only features of their bottom-up input, but generate ‘predictions’ or hypotheses about the likely causes of their sensory inputs, and match these to the incoming sensory input to compute the mismatch or ‘prediction error’ (Friston, 2005; Kok, 2016; Rao & Ballard, 1999; Stephan et al., 2019).

The basic assumption of such coding schemes is that predictions are conveyed via feedback connections originating in higher levels of the cortical hierarchy and are compared with lower level inputs to calculate the extent of prediction error relative to concurrent incoming input. Any ensuing error signal is relayed back up the hierarchy via feedforward connections to update predictions in higher level regions and minimise prediction error at each level of the hierarchy, until all errors are reconciled (Bastos et al., 2012; Friston, 2005; Kanai et al., 2015). A key feature of this account is that these sub-populations of neurons are proposed to reside in different cortical layers. For

instance, predictions conveyed through feedback connections originate from the deep layers of higher-order regions and are passed to the agranular layers of lower-level regions, whereas prediction error signals are passed up via feedforward connections and thought to reside in lower level superficial layers sending input to the middle layers of higher level areas (Bastos et al., 2012; Friston, 2005; see Fig 1.3.C).

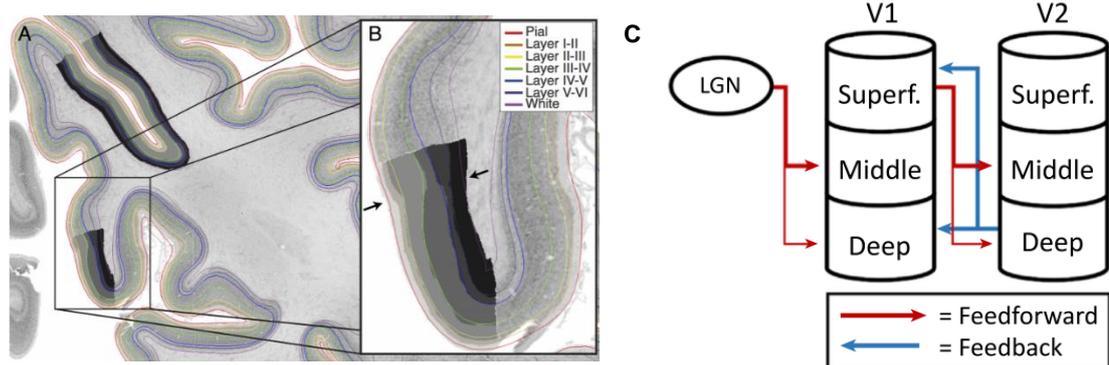


Figure 1.3. (A) Depiction of the six cortical layers intersected onto a section of occipital cortex. (B) Arrows denote the boundary between V1 and V2 (Secondary Visual Cortex). Adapted from Wagstyl et al. (2020). (C) Schematic representation of proposed extrinsic feedforward (red) and feedback (blue) connections across layers, between the LGN, V1 and V2; taken from Lawrence, Formisano, et al. (2019).

There have been a range of experiments designed to provide empirical support for such theoretical assumptions, and studies examining the laminar activity profile of expectation, in particular, may be able to provide insight into a potential resolution to conflicting neural prediction findings. As such, it is possible that contrasting upweighting (e.g., sharpening) and downweighting (e.g., expectation suppression) neural effects observed across literatures could be explained by separate populations of neurons residing in different layers of the cortex that encode different aspects of expectation. For instance, neural populations in superficial layers proposed to encode prediction error may be able to account for increased unexpected activity in expectation suppression findings, while populations in deep layers thought to encode hypotheses or

predictions about incoming sensory input may display effects more akin to sharpening, that highlight expected representations. With modern techniques now allowing examination of activity patterns across different cortical layers in human sensory cortices, research is more equipped to answer this question, and such a possibility is considered in the following section.

1.4.2. Evidence for distinct prediction and prediction error populations

Support for the central tenet of predictive coding theories that prediction (P) and prediction error (PE) populations are distinctly represented within a cortical region has gained empirical support from invasive single neuron recording studies in monkeys. For instance, Bell and colleagues (2016) manipulated the probability of stimulus categories (e.g., faces and fruits) shown to macaque monkeys while measuring neural responses in inferior temporal cortex (IT). They found that neural firing rates were higher for unexpected stimuli relative to expected stimuli, while multivariate analyses revealed that stimuli that were expected could be decoded with greater accuracy. While neuronal cortical depth was not recorded here, the authors examined the correlations across neuronal populations that were selective to either predicted or surprising stimuli, and revealed that there was no correlation between those populations. Such evidence might provide preliminary support for the claim that there are separable P and PE populations within sensory cortical regions such as IT, perhaps subserving distinct mechanisms.

Other studies using similar methods but probing different cortical layers have adopted figure ground segregation tasks in macaques to investigate the laminar profiles of feedforward and feedback processes in visual cortices primarily. For example, Self and colleagues (2013) report that the strongest effects of top-down figure filling-in were present in superficial and deep layers of V1, avoiding the middle input layers (Self et al., 2013). Additionally, when monkeys are presented with Kanizsa figures that allude

to shape contours, neural responses to the illusory edges are found to be strongest in superficial and deep layers of V1 and V2, and are interestingly present earlier in V2 (Lee & Nguyen, 2001). These findings support the notion that feedback connections from higher areas represent predictive information about stimulus features on the basis of prior knowledge, and provide some evidence for the idea that top-down prediction processes operate via connections to the agranular layers of lower-level sensory regions.

While many of these studies examined monkeys due to the invasive nature of laminar electrodes, the relatively recent development of ultra-high resolution 7-Tesla (7T) MRI methods has allowed examination of human laminar activation patterns. This high resolution technique means that cortical voxel sizes of roughly 0.8 mm can be acquired, where studies can measure laminar activity profiles by dividing voxels in the cortical surface into three grey matter voxel bins (representing ‘deep’, ‘middle’ and ‘superficial’ layers). This method has grown widely popular in the last decade, with studies largely interested in examining the evidence for separable top-down and bottom-up influences on perceptual processing in human sensory cortices. For instance, investigation into top-down influences such as attention and working memory on perception have provided evidence for the claim that feedback targets the deep and superficial layers of lower-level (e.g., V1) regions, similarly to in monkeys. Namely, visual features that are attended (Lawrence, Norris, et al., 2019) and held in working memory (Lawrence et al., 2018) generate a stronger neural response in agranular layers of V1, while orthogonal bottom-up manipulations were found to activate all layers but most strongly middle layers. Notably, these studies are consistent with the established functional neuroanatomy of feedback and feedforward projections, terminating in deep and superficial, or middle layers, respectively (Anderson & Martin, 2009; Felleman & Van Essen, 1991; Rockland & Virga, 1989).

Furthermore, a study by Kok and colleagues (2016) presented Kanizsa figures to human participants and examined top-down influences on laminar activity patterns resulting from illusory triangular edges in the absence of bottom-up input. Their findings revealed selective activation of only the deep layers of V1 in response to the illusory contours, while bottom-up input that did not induce illusory figures activated all cortical layers (Kok, Bains, Van Mourik, et al., 2016). Some of these authors have also demonstrated increased activity in deep layers for expected but subsequently omitted visual stimuli that is distinct from patterns evoked by presented expected stimuli that activated all layers equally (Aitken et al., 2020). However, it is worth noting that another study similarly examining top-down prediction influences in the absence of bottom-up input has alternatively demonstrated, using MVPA, that the content of occluded visual scenes could be decoded only in the most superficial cortical layer 1 of V1 (Muckli et al., 2015). Differences between these findings have been suggested to reflect differences in analytical approaches, which may relate to known limitations of the BOLD signal across layers, and is a possibility that will be discussed in more detail in Chapter 5. Nevertheless, these findings cleverly demonstrate top-down predictive influences in the absence of bottom-up input and suggest that higher-order predictions conveyed through feedback connections in humans indeed likely target the agranular layers of V1. Although this evidence supports assumptions that there is a sub-set of neuronal populations encoding predictions in deep layers, there is a notable lack of evidence demonstrating how perceptual processes associated with unexpected events operate and whether these map on to PE encoding in superficial layer populations.

1.4.3. Open questions

A possible reconciliation to the disparity between neural findings of expectation suppression and sharpening across domains could come from claims that there are two

distinct sub-population of neurons representing P and PE processing across different layers of cortical regions. In particular, it is possible that increased representation of expected outcomes in deep cortical layers (encoding predictions) could account for sharpening findings, while increased representation in superficial layers (encoding prediction error) of unexpected events could account for findings of expectation suppression. To investigate this possibility, the final experiment of this thesis uses 7T fMRI to examine laminar activity profiles in human V1 and V2 associated with expected and unexpected action-outcomes (Chapter 5).

1.5. Thesis outline

Chapters 2 and 3 investigate the claim that action predictions influence touch perception in a distinct fashion from other types of prediction in broader sensory cognition.

Chapter 2 adopts the typical force judgement paradigm used in the action literature to demonstrate tactile attenuation and adds an additional condition that does not require active contact. This is in order to examine a potential confounding influence in these typical set-ups (Experiment 1). Chapter 3 presents three experiments that combine a version of this force judgement paradigm with manipulations of prediction according to learned statistical probabilities, typical in sensory cognition literature. In these experiments, participants are trained on the association between their action and the resulting tactile force and comparisons are now made between force judgements for expected and unexpected outcomes (Experiments 2-4). Computational modelling examines whether expectation effects could be explained by perceptual biases or alternatively by response biases. Together, these experiments aim to determine whether action predictions exert a distinct attenuating influence on touch perception or instead demonstrate upweighting influences – comparable to in sensory cognition domains that similarly manipulate predictability.

Chapter 4 investigates the neural nature of prediction mechanisms in V1 by manipulating the relationship between action and visual grating orientations, again making comparisons between expected and unexpected events. Experiment 5 reports an orientation discrimination task that examines how expectations about stimulus orientation influence perceptual discrimination responses, and Experiment 6 reports an fMRI experiment investigating how action predictions are represented in patterns of V1 activity evoked by expected and unexpected stimuli. These experiments aim to determine whether neural prediction mechanisms operate in a distinct fashion during action relative to those observed in the broader sensory cognition literature.

Finally, Chapter 5 investigates how differences in action-outcome contingencies influence perception and uses these findings to inform a paradigm examining the laminar activity profiles of expected and unexpected events in visual cortical regions. Experiment 7 therefore reports a psychophysical study looking at the influence of degraded action-outcome associations on perceptual discrimination judgements, while Experiment 8 reports a 7T fMRI experiment that aims to investigate how expected and unexpected events are represented across cortical layers of visual regions. The primary aim of these experiments is to determine whether discrepant neural findings, demonstrating both sharpening and expectation suppression, could be reconciled by a conclusion from predictive coding theories that P and PE are represented by distinct neural populations residing in different layers of cortical regions.

Chapter 2: Examining tactile attenuation in action

Chapter 1 therefore highlights a disparity between two theories that put forward how prediction influences perception. Cancellation theories propose that predictions about the expected sensory consequences of action attenuate perception in order to prioritise processing of unexpected outcomes that drive corrective action (Press et al., 2020b; Wolpert & Flanagan, 2001). On the other hand, theories of how prediction influences perception outside of action propose that predictions about what we expect to be there – based on the common statistical regularities of the world – are combined with incoming sensory input to bias our perceptual experiences *towards* what we expect (de Lange et al., 2018; Kersten et al., 2004; Knill et al., 1996). A notable difference between studies that support each of these theories across domains is that the majority of literature supporting cancellation accounts examines tactile perception while the literature supporting Bayesian accounts largely examines visual perception. Thus, it is possible that the paradox could be reconciled by a conclusion that sensorimotor predictions influence *tactile* perception differently from other types of prediction. For example, suggestions that action predictions shape tactile perception in a qualitatively distinct way include proposals that differences relate to tactile events being body-related (Dogge et al., 2019) and tightly coupled with the motor system (Kusnir et al., 2019), in a way that many predicted visual or auditory events are not. However, given that the studies examining tactile perception manipulate predictability in a different way from visual studies, such a conclusion would be premature. This chapter aims to examine the nature of tactile attenuation reported in the action literature while controlling for some potential methodological confounds in typical set-ups supporting cancellation theories.

2.1. Experiment 1

Experiment 1 employs an adaptation of typical force judgement paradigms popular in the action literature for demonstrating tactile attenuation. Such paradigms (e.g., Bays et al., 2006, 2005; Kilteni, Houborg, & Ehrsson, 2019; Wolpe et al., 2018) require participants to move an active right index finger towards a static left finger which receives stimulation mechanically. Tactile stimulation on the left finger is reported as less forceful when accompanied by action (e.g., in ‘active’ conditions), relative to the same forces but when no action is performed, or when the active finger remains still (in ‘passive’ conditions).

A key claim in some of these studies is that predictions about the resulting stimulation can only be formed when the active finger makes contact during movement (e.g., with a force generator in these paradigms), and this contact is therefore necessary for observing predictive tactile attenuation effects (Bays et al., 2006; Kilteni & Ehrsson, 2017a). However, it is unclear what type of adaptive mechanism might depend on contact during action to form predictions about the resulting sensory consequences, because actions that do not result in contact can still predict sensory outcomes if underlying predictive mechanisms are defined according to stimulus probabilities (de Lange et al., 2018). It is likely that we have learned over the course of our lives that actions frequently result in contact e.g., to touch your own hand requires that the active hand will receive stimulation as well as the receiving hand, and therefore we might expect certain sensory outcomes on the basis of active contact. This experience is however not ubiquitous, and we are also likely to experience sensory outcomes resulting from contactless actions e.g., automatic hand driers result in warm air being blown on to our hands as a sensory outcome of action which does not require active contact with any objects. Therefore, if sensorimotor prediction is underpinned by mechanisms defined

according to learned associations, then active-finger contact should not be required to form such predictions *per se* and should in principle operate similarly in no active contact situations.

Additionally, downweighting mechanisms that are exclusively produced by actions that result in contact could be explained by a number of non-predictive factors known to influence perception during action (Firestone & Scholl, 2016; Press & Cook, 2015). Specifically, typical studies demonstrating tactile attenuation during action compare the perception of events in the presence or absence of action, or when events are coincident versus delayed with respect to action (Bays et al., 2005, 2006; Blakemore et al., 1998; Kiltner et al., 2019; Shergill et al., 2013; Wolpe et al., 2018). In these experiments, it is assumed that the sensory events which coincide with action are more predicted, explaining why perception of them is attenuated. However, tactile attenuation in active conditions could in principle be generated by generalised sensory gating mechanisms when there is active contact relative to passive conditions. It is widely acknowledged that performing actions generates identity-general gating of sensation on a moving effector – with studies demonstrating greater reduction in detectability of tactile stimulation during movement at stimulation sites closest to the moving finger (Williams et al., 1998; Williams & Chapman, 2000). These effects are thought to be mediated by neural mechanisms observed at the earliest relay in the spinal cord (Seki & Fetz, 2012). Identity-general gating effects can also be observed in a variety of other sensory systems, for example, saccadic suppression is thought to minimise transient blur during a visual gaze shift, and sensory filtration has been found to suppress primary auditory cortex activity associated with self-vocalisation, suggested to allow attunement to externally produced sounds (Crapse & Sommer, 2008). In humans, these mechanisms are also not thought to be predictive, because they influence perception of events

regardless of whether they were predicted outcomes of action e.g., gaps in ongoing wrist vibrations are not a typical outcome predicted by juggling but are perceptually attenuated during this task (Juravle & Spence, 2011).

Force judgement paradigms have been employed by action researchers partly to allow for dissociation of identity-specific prediction effects from those of identity-general gating effects by probing perception of events presented to static effectors rather than active ones (Bays et al., 2006). However, there are a number of ways that perception of concurrent gated stimulation on an active finger – as generated when pressing a device in such paradigms – could still bias responses about stimulation on the static left finger. One possibility relates to response biasing (Firestone & Scholl, 2016), such that participants might be biased towards responding about the gated sensation when it appears synchronous with passive sensation. Another possibility is that perception itself is biased towards the gated sensation when the two events occur in close proximity, becoming spatiotemporally bound such that the tactile stimulation presented to the left index finger represents an average of the forces experienced by both active and passive effectors. Binding of signals in close spatiotemporal proximity may be adaptive to portray to the brain that the two signals have come from a common source (Conrad et al., 2012). However, such explanations are speculative and it is unclear how active contact might bias passive sensation during active contact. Nevertheless, the variety of possible explanations suggests that such manipulations might still confound predictive attenuation with sensory gating influences, relative to when such gated active stimulation is not present during action.

Experiment 1 therefore presents an adaptation of the typical force judgement paradigm to examine this possibility by comparing effects in two conditions. A ‘Contact’ condition closely resembled that employed by previous work, whereby participants

moved their right index finger ('active finger') to make contact with a button, generating a mechanical force via a solenoid to their left static index finger below (see Fig 2.1.A). The 'No Contact' condition was included to see whether similar effects would be found when an identical downward motion of the active finger triggered the same left-hand finger stimulation, but where the active finger did not make contact with the button. Stimulation to the left hand was instead triggered via infrared motion-tracker detection of right-hand finger motion. This condition was included to investigate the nature of attenuation effects when possible confounds associated with active contact are reduced, because if sensorimotor prediction is underpinned by predictive mechanisms defined according to event probabilities, then active-finger contact should not be a requirement to form such predictions, though it may still influence perception. Participants made force intensity judgements about the mechanical stimulation when accompanied by action ('active' trials), or when all fingers were inactive ('passive' trials).

The Contact condition was hypothesised to replicate the findings of Bays et al. (2005; 2006), such that participants would perceive tactile events to be less forceful when accompanied by action compared to passive trials. However, if the operation of other mechanisms contributes to attenuation in Contact conditions, the pattern should be different in the No Contact condition when this influence is removed. For instance, events in the No Contact condition may possibly even be perceived with greater intensity when accompanied by action, as would be predicted by upweighting accounts.

2.1.1. Method

Participants

Thirty participants (16 female, mean age = 25.53 years [$SD = 5.25$]) were recruited from Birkbeck, University of London and paid a small honorarium for their participation. Eight participants were replacements for those who could not complete the perceptual discrimination and an additional four for exclusions due to technical error. Three participants had PSEs in at least one of the four conditions that fell marginally outside the range of presented stimuli (range of stimuli = 1-7, range of PSEs = 0.3-7.7). Visual inspection determined that psychometric function fits were reasonable so these participants were included in the sample – importantly, the key pattern of results was the same when they were also excluded. All participants reported no current neurological or psychiatric illness and provided written informed consent prior to participation. The sample size was determined a priori on the basis of pilot testing to estimate effect size – to have at least 80% power of detecting medium effect sizes ($d = 0.5$) – and parametric assumptions were met. The experiment was performed with local ethical committee approval and in accordance with the ethical standards laid down in the 1964 Declaration of Helsinki for both experiments.

Tactile stimuli

Tactile forces were mechanically delivered to participants' left index finger via a solenoid and were presented via a soundcard. In order to present tactile forces that could be perceived as a single tap, a non-sinusoidal half-square wave of low frequency (16.67 Hz) was presented for 30 ms for each individual force. Stimulus intensity was manipulated via the amplitude of the square wave and seven logarithmically spaced

amplitudes were selected to provide a range of detectable intensities. The fourth (middle) intensity was used as the reference force.

Procedure

The experiment was conducted in MATLAB using the Cogent toolbox (<http://www.vislab.ucl.ac.uk/cogent.php>). Participants held their left-hand palm upwards with their index finger positioned against a solenoid (diameter of metal rod = 4 mm; diameter of solenoid = 15 mm; TACT-CONTR-TR2, Heijo Research Electronics) sitting on the apex of the fingertip. Their right-hand rested on a right armrest, positioned such that the index finger distal phalange was directly above the left-hand distal phalange, but rotated 90 degrees anticlockwise relative to their left hand (see Fig. 2.1.A). A small button box or infrared motion tracker (Leap Motion Controller using the Matleap MATLAB interface) was placed on a shelf supporting the solenoid (depending on block type, Contact or No Contact respectively), positioned directly above it.

At the start of each trial, participants were cued onscreen to move their right index finger ('move'; Active trials – 50%) or remain stationary ('do not move'; Passive trials – 50%). In Contact blocks, participants' right hand was positioned 5 cm above their left hand, and in No Contact blocks it was moved to 12 cm above to allow movements to be made without touching the lower shelf holding the motion tracker. In both Contact and No Contact blocks participants were required to hold their right-hand palm facedown, with their arm parallel with the computer monitor and the coronal body midline. The experimenter provided demonstration of appropriate action execution for both conditions before the onset of the experiment. These steps ensured that participants' starting position remained the same regardless of trial type (active vs passive), and that

actions were executed in the same way across blocks (Contact, No Contact; bar the extent of contact at termination).

On Active trials, they rotated their index finger downwards at the metacarpophalangeal joint. The target stimulus was delivered to the left index finger for 30 ms immediately after motion was detected by a button press in Contact blocks, or in No Contact blocks as soon as the active finger movement achieved approximately equivalent movement distance to Contact blocks (distance of at least 20 mm from the stationary finger starting position at trial onset). Piloting confirmed that stimulation was in apparent synchrony with movement termination in both Contact and No Contact blocks.

After 1000 ms, a reference stimulus was presented for 30 ms. The target stimulus presented one of seven logarithmically-spaced supra-threshold forces, and the reference stimulus always presented the fourth (middle) force. After an inter-stimulus interval (ISI) of 300 – 500 ms, participants were asked which tap was more forceful, responding with a left foot pedal for the first stimulus and a right foot pedal for the second stimulus. The next trial started after an inter-trial interval (ITI) of 1000 ms. In Passive trials, the target stimulus was delivered 500 ms after the cue to remain still. Due to the fact that the Contact vs No Contact manipulation only affected active trials, passive trials were identical in both Contact and No Contact blocks and served as comparisons within blocks. In all conditions, participants' left-hand was positioned laterally from the body midline and in line with the shoulder. Participants' hands were visually occluded during the experiment and white noise was played through headphones (53 Db; piloting confirmed that this level resulted in inaudible solenoid movement) throughout testing.

There were 560 trials in total; 140 for each of the Active and Passive conditions, and in each of the Contact and No Contact blocks. The order of blocks was counterbalanced

across participants and trial type order was randomised. Participants completed eight practice trials before the main test blocks.

Modelling psychometric functions

Participant responses were modelled by cumulative Gaussians to estimate psychometric functions, using the Palamedes Toolbox (Prins & Kingdom, 2018) in MATLAB. This procedure was performed separately for Active and Passive trials during the test phase. Points of Subjective Equivalence (PSEs) were calculated from the psychometric functions, which describe the point at which participants judge the target and reference events to have equal force. Lower values are indicative of more intense target percepts.

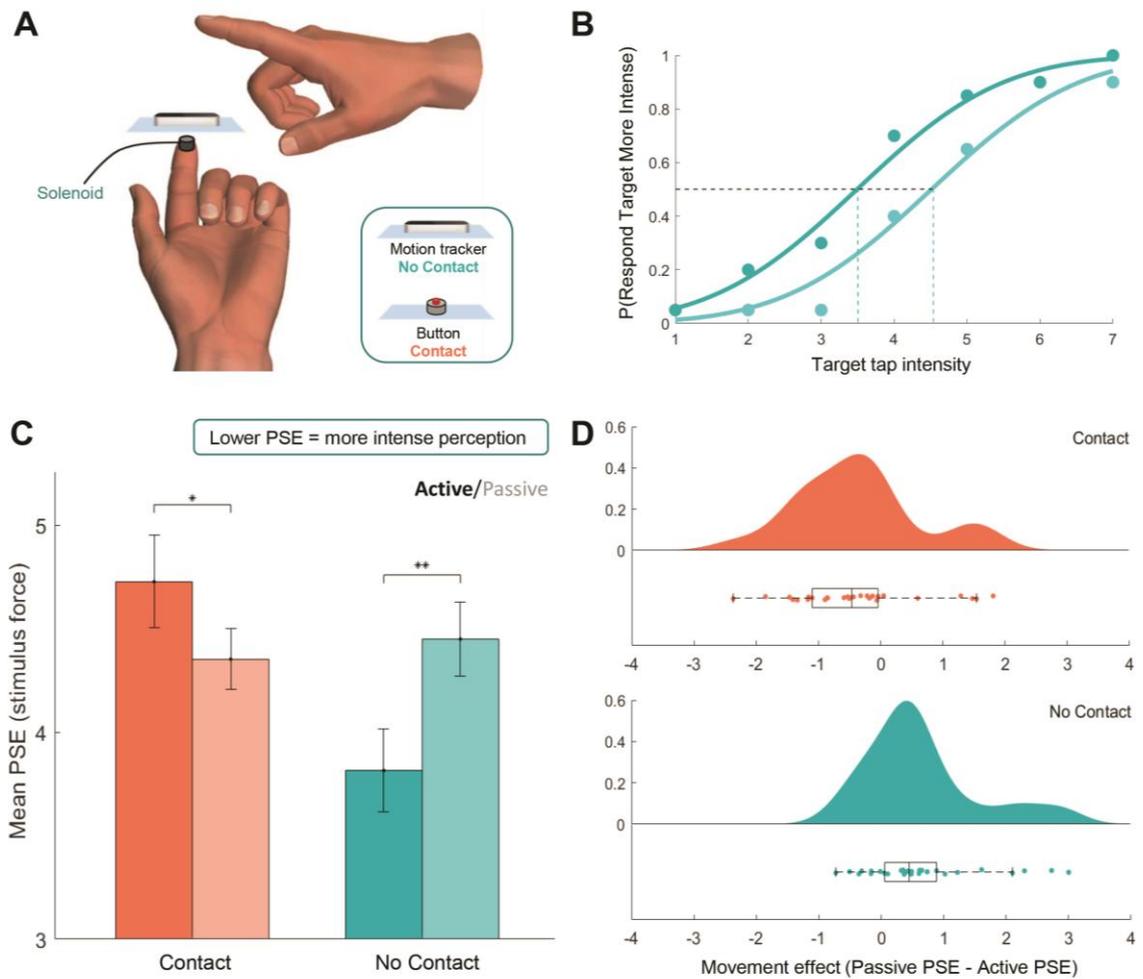


Figure 2.1. (A) On each trial, participants made downward movements with their right index finger either over a motion tracker (No Contact condition), or towards a button with which they made contact (Contact condition). Each movement elicited a tactile punctate event to the left index finger positioned directly below. (B) PSEs were calculated for each participant and this data represents an example participant in the No Contact condition for Active [dark blue] and Passive [light blue] trials. (C) Mean PSEs (\pm SEM) were higher in Active than Passive trials in the Contact conditions, but lower in Active than Passive trials in the No Contact condition. Lower PSEs indicate more intense target percepts (* $p < .05$, ** $p = .001$). (D) PSE effect of movement (Passive – Active) for the Contact (top) and No Contact (bottom) condition, plotted with raincloud plots displaying probability density estimates (upper) and box and scatter plots (lower). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive effects of movement indicate more intensely perceived active target events relative to passive events, but negative values indicate the reverse – less intensely perceived active events.

2.1.2. Results and Discussion

PSE values were analysed in a 2x2 within participants ANOVA, revealing no main effect of Contact ($F(1, 29) = 3.11, p = .089, \eta_p^2 = .10$) or Movement ($F(1, 29) = 1.24, p = .274, \eta_p^2 = .04$). However, there was a significant interaction between Contact and Movement ($F(1, 29) = 15.39, p < .001, \eta_p^2 = .35$), driven by higher PSEs in Active ($M = 4.73, SD = 1.22$) compared to Passive trials ($M = 4.35, SD = .80$) in the Contact condition ($t(29) = 2.07, p = .047, d = .38$), but lower PSEs in Active ($M = 3.82, SD = 1.11$) compared to Passive trials ($M = 4.45, SD = .97$) in the No Contact condition ($t(29) = -3.80, p = .001, d = .69$). Thus, when an active finger receives cutaneous stimulation – through contact with the button – simultaneous with stimulation to the left-hand effector, tactile events are perceived less intensely during action than when the active finger stays still. Conversely, when the active finger does not receive such stimulation in No Contact conditions, tactile events are perceived *more* intensely during action.

2.2. General Discussion

Experiment 1 revealed that when a moving right-hand finger makes contact with a button simultaneously with left index finger stimulation, tactile events are perceived less intensely during movement, yet are perceived more intensely during movement when active contact does not occur. These results replicate previous tactile attenuation findings using force judgement paradigms (Bays et al., 2005; Kilteni et al., 2019), but only when the specifics of the paradigm were replicated such that the active finger makes contact with a button.

The No Contact findings may however be seen as consistent with assumptions that active contact is necessary requirement to observe tactile attenuation, and thereby to form predictions about the resulting sensory consequences, however it is unclear why a

flip in effects would be observed when this feature is removed. Instead, this finding raises the possibility that tactile attenuation in active contact set-ups could be explained by a number of possible non-predictive influences of action on perception (Press & Cook, 2015), which potentially mask predictive influences relative to when contact does not occur.

It has already been considered that attenuation in force judgement paradigms may in fact be generated by identity-general gating mechanisms when comparisons are drawn between active and passive conditions. For example, spinal gating mechanisms – which suppress tactile sensitivity on a moving finger (Williams et al., 1998; Williams & Chapman, 2000), irrespective of what is predicted – could contribute to previous reports of attenuation in the force judgement paradigm if attenuated percepts on the active finger bias reports about forces on the static finger. For example, in a study by Shergill, Bays, Frith and Wolpert (2003), participants were required to match the perceived force of a previous externally generated tap on their left finger with an action from their own right finger. Here, participants tended to overestimate the force required from their right finger, interpreted by the authors to suggest that touch perception was suppressed on the passive finger due to prediction (Shergill et al., 2003). This overestimation could also be explained by non-predictive identity-general mechanisms on the active finger, requiring them to apply additional force to counteract this suppression during action. Comparison conditions in action studies also frequently compare perception of ‘predicted’ action consequences against perception where the action and sensory events are temporally misaligned (Bays et al., 2006; Blakemore et al., 1999; Kiltner et al., 2019; Shergill et al., 2013). These types of manipulation could arguably similarly yield large influences of identity-general gating when the overlap between action and sensory events is great, relative to when the sensory event occurs at a delay following action.

The influence of sensory gating on active effectors is thought to be removed in typical force judgement paradigms because they measure forces on a passive effector during movement (Bays et al., 2005, 2006). However, the present results might indicate that such influences still remain. Bays and colleagues (2006) designed an experiment to investigate the specific influence of sensory gating during active contact in these force judgement paradigms, but with results that contrast the No Contact findings of Experiment 1. In their experiment, two groups of participants were employed to perform the typical force judgement task. In a group A, akin to the Contact condition in Experiment 1, participants frequently made active movements resulting in contact with the force sensor, and received stimulation to their left static index finger simultaneous with action. On infrequent occasions and unknown to the participant, the force sensor was moved such that there was no active contact when participants made actions towards the force sensor. The authors report effects of tactile attenuation in both types of trials regardless of contact, and claim that these effects resulted from predictions that could be formed on the basis of the contact they expected to make during action. Furthermore, in a group B, akin to Experiment 1's No Contact condition, participants performed the same actions and never made contact with the sensor. This time, tactile attenuation effects are abolished, which is suggested to result from no expectations regarding making active contact. This combination of findings is thought to provide evidence for predictive downweighting mechanisms during active contact because attenuation is removed when such contact is not expected (Bays et al., 2006).

It is noticeable that when active contact is removed in group B of Bays et al.'s (2006) study (similarly to the No Contact condition), they observe null effects, while an enhancement effect is observed in the comparable condition of Experiment 1. There are a number of possible explanations for this difference, with one possibility being that

Bays et al.'s (2006) study lacked sufficient power to detect a significant flip over to enhancement when only 10 participants were tested in comparison to a sample three times the size in Experiment 1. Such a flip might be expected on the basis of the current findings, if enhancement effects in no contact conditions reflect predictive upweighting mechanisms when influences of sensory gating are controlled. However, there are a variety of influences and methodological differences that could explain the differences between these conditions. As well as contributions of gating as speculated earlier, there may also be differences in the executive task demands, especially when active and passive comparisons are made. For instance, it has been demonstrated that perceptual sensitivity for detecting tactile events presented across different limbs is reduced when participants are also engaged in a secondary motor task (Gallace et al., 2010). Further, classic accounts of working memory would predict fewer resources for processing target stimuli in the presence of additional tasks such as action (Baddeley, 1996; Press et al., 2020a), relative to passive judgements. Therefore, further tests will be necessary to determine the precise mechanism accountable for the flip observed in the present experiments when active contact was manipulated, as well as to determine the predictive nature of the enhancement effects observed.

It is worth noting that there was one other difference between the Contact and No Contact conditions reported in Experiment 1. The Contact condition aimed to replicate the set-up of the Bays et al. (2005; 2006) studies as closely as possible, yet introducing an additional No Contact condition rendered the hands at a separation distance across conditions that could not be precisely matched due to the infrared movement tracking. In principle, this greater hand separation could have interacted with the operation of any predictive mechanism, considering that studies have demonstrated reduced attenuation the further the hands are separated (Kilteni & Ehrsson, 2017a). It is also possible that

closer spatial proximity of the left and right hands in Contact conditions was more likely to render a representation binding both sensations due to proximity (Conrad et al., 2012) relative to when hands are further separated in the No Contact condition.

However, it is unclear why any predictive attenuation theoretical accounts would predict that this small increase in hand separation in the No Contact condition to induce a flip to enhancement effects other than for the explanation already hypothesised. Importantly, all conclusions relate to the simple effects within the Contact and No Contact conditions, so this additional difference should not alter the conclusions.

Regardless of the precise mechanisms generating the perceptual effects across conditions and studies, the important conclusion from Experiment 1 is that the typical tactile attenuation effect was replicated when the methods were most precisely replicated, but a small adjustment to the method reversed the direction of effects in line with upweighting influences on perception. It could be suggested from this finding that active contact potentially masks underlying predictive processes when perceptual comparisons are made between active and passive conditions – rendering effects that are possibly reflective of non-predictive influences of action on perception. Yet, although this enhancement effect observed is consistent with findings in visual perception (Yon et al., 2020; Yon & Press, 2017, 2018) and in line with upweighting theories of prediction, it is impossible to draw such conclusions before manipulating statistical probabilities between actions and sensory outcomes unconfounded by non-predictive influences on perception.

Chapter 3: Does expectation enhance or attenuate tactile perception?

Experiment 1 replicated typical tactile attenuation effects in a force judgement paradigm, but only when specifics of the paradigm were replicated such that the active finger makes contact with a button. Interestingly, tactile events produced by actions that do not result in active contact – and therefore where gating of cutaneous stimulation on the active effector cannot determine effects – were rated as *more* forceful than passive events. This enhancement effect may be plausibly explained by upweighting mechanisms that propose that sensory input is combined with expectations about the statistical regularities of our sensory world to bias perception towards what we expect (de Lange et al., 2018; Kersten et al., 2004). This interpretation would also be consistent with some recent evidence in the visual domain demonstrating upweighting influences of action on *visual* perception (e.g., Yon et al., 2018, 2020; Yon & Press, 2017). However, it would be premature to conclude that any form of prediction mechanism is underlying influences of action on perception before manipulating the predictability of tactile events according to their conditional probabilities – akin to in studies in the wider sensory cognition literature.

The defining feature of prediction mechanisms is that they operate according to stimulus probabilities (de Lange et al., 2018). As such, effects of prediction mechanisms are typically measured by presenting events with high and low conditional probabilities, allowing comparison of perception of ‘expected’ (e.g., 80% likely based upon a preceding cue) and ‘unexpected’ events (e.g., 20% likely; Cheadle et al., 2015; Han et al., 2019; Kok et al., 2012, 2013, 2014; Kok & Turk-Browne, 2018; Richter & de Lange, 2019). In contrast, typical experiments demonstrating tactile attenuation during action compare the perception of events in the presence or absence of action, or when

events are coincident versus delayed with respect to action (Bays et al., 2005, 2006; Blakemore et al., 1998; Kilteni et al., 2019; Kilteni & Ehrsson, 2017a; Shergill et al., 2013; Wolpe et al., 2016, 2018). In these experiments, it is assumed that the sensory events that coincide with action are the (more) predicted consequences, which explains why perception of them is attenuated. However, without manipulating conditional probabilities it is unclear how these effects compare to predictive effects outside of action, and specifically whether any of these effects are indeed reflective of predictive mechanisms, considering there are a range of non-predictive mechanisms known to influence perception during action (Press et al., 2020a; Press & Cook, 2015).

Therefore, in order to address the main theoretical question of this thesis regarding whether prediction influences perception distinctly in action domains, Experiments 2-4 isolated the particular functional influence of prediction by manipulating conditional probabilities between actions and outcomes. Perception of tactile events was now compared when events were expected or unexpected, always in the presence of action – to overcome confounds associated with active vs passive comparisons. These experiments tested whether tactile events predicted by action are perceived less intensely than unexpected events – as downweighting accounts would propose – or instead are perceived as more intensely to aid veridicality – consistent with upweighting theories of prediction outside of action.

3.1. Experiment 2

Experiment 2 therefore examined the predictive influence of action on tactile perception using a statistical learning paradigm, where participants performed one of two movements that were predictive of one of two tactile effects. Comparing expected and unexpected conditions always in the presence of action – rather than the typical active vs passive comparison – allows isolation of predictive influences on tactile perception, while controlling for possible non-predictive influences on perception during action that are not comparable in passive conditions (Press & Cook, 2015). Additionally, presenting two action types and two stimulation types enables comparison of perception of events that were expected or unexpected while controlling for repetition effects (see section 3.5. for discussion).

Specifically, participants' actions required rotating their right index finger either upwards or downwards at the metacarpophalangeal joint and over a motion tracker. Such movement triggered the delivery of a target tactile stimulus to either the left index finger or left middle finger, presenting tactile events during action in the absence of active contact, similar to the No Contact set-up in Experiment 1. During a training session, an action (e.g., upward movement) was perfectly predictive of the tactile event (e.g., stimulation to the left index finger). Participants were presented with a test session 24 hours later, in which the action-outcome relationship was degraded to additionally measure perception of unexpected events whereby the expected finger was stimulated on 66.7% of trials, and the unexpected was stimulated on the remaining 33.3% of trials. Training and test sessions were presented on consecutive days to allow statistical learning during training to be consolidated with sleep – a process that has been widely demonstrated to improve effects of learning (Durrant et al., 2011; Gagnepain et al., 2012). While downweighting accounts predict that expected tactile events will be rated

less intensely (forceful) than unexpected events, upweighting theories predict that expected events will be rated more intensely than unexpected events.

3.1.1. Methods

Participants

Thirty participants (20 female, mean age = 22.80 years [$SD = 3.18$]) were recruited in the same way as in Experiment 1. Six participants were replacements for those where acceptable psychometric functions could not be modelled to their responses, where they were unable to follow instructions concerning movement performance, or where there was technical error. One participant's PSE scores were winsorized to meet the normality assumptions of parametric tests (from $z = 3.34$ to $z = 3$; Tukey, 1962). The sample size was matched to power estimates from Experiment 1 and is the same for all experiments in this chapter. Ethical approval was also obtained in the same way as Experiment 1 and applies similarly across the chapter.

Procedure

The following changes were made relative to Experiment 1. Participants were positioned with their left index and middle fingers making contact with independent solenoids (Fig. 3.1.A). At the beginning of each trial an arbitrary cue (either a square or circle) instructed participants to move their right index finger either upwards or downwards from the metacarpophalangeal joint, tracked by an infrared motion sensor. When motion was detected, the target stimulus was delivered to either the left index or middle finger for 30 ms, resulting in apparent synchrony of stimulation with movement. All other stimulus presentations and timings were the same as in Experiment 1. During training, participants' right-hand index-finger action (e.g., downwards movement) was 100% predictive of the location of left-hand tactile events (e.g., index finger). In a test

session 24 hours later, the action-outcome relationship was degraded to measure perception of expected and unexpected events – the expected finger was stimulated on 66.7% of trials, and the unexpected was stimulated on the remaining 33.3% of trials.

It should be noted in explicating the logic of the procedure here that any action predictions should determine where stimulation will be received, rather than its intensity. However, in Bayesian models it is assumed that enhanced detection and intensity of expected events relates to the precision of the estimate (Brown et al., 2013; e.g., a force precisely estimated to have occurred on a certain region of tactile space should be felt as more intense because of the precise estimate of spatial information, rather than an estimate of the force per se). Since these models assume that predictions enhance the precision of resultant estimates, they would also predict enhancements in perceived force (and indeed other sensory attributes, like brightness or loudness).

There were 420 trials in each session. Trial order was randomised and the action-stimulus mapping was counterbalanced across participants. The cue-action mapping was switched half way through the main trials and this order was counterbalanced across participants. Participants completed 12 practice trials before the main session trials.

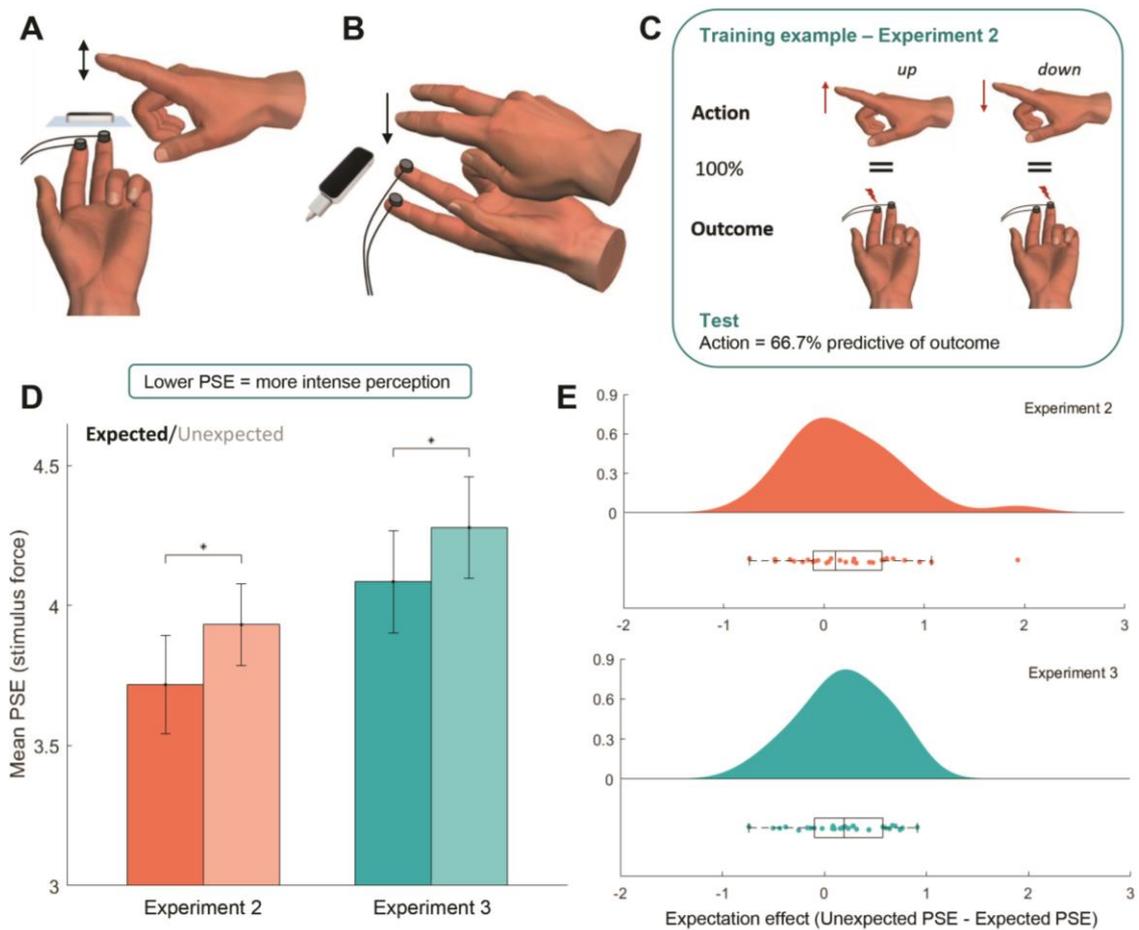


Figure 3.1. (A) Experiment 2 set-up. Participants made a downwards or upwards movement with their right index finger over a motion tracker, eliciting punctate tactile events to the left index or middle finger. (B) Experiment 3 set-up. Participants made only downwards movements, now with either their right index or middle finger. (C) Example training procedure for Experiment 2. Movements were perfectly predictive of tactile events in the training session, and 66.7% predictive in the test session. The same procedure was adopted in Experiment 3 but with different action types. Flash symbols illustrate location of stimulation. (D) Mean PSEs (\pm SEM) were lower for expected than unexpected trials in both Experiment 2 and Experiment 3, indicating more intensely perceived target stimuli ($* p < .05$). (E) PSE expectation effect (Unexpected – Expected) plotted with raincloud plots (Allen et al., 2019) displaying probability density estimates (upper) and box and scatter plots (lower), for Experiment 2 (top) and Experiment 3 (bottom). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant ($N=30$). Positive effect values indicate more intensely perceived expected events.

3.1.2. Results and Discussion

PSE values were lower on expected trials ($M = 3.72$, $SD = .96$) than unexpected trials ($M = 3.93$, $SD = .80$; $t(29) = -2.12$, $p = .041$, $d = .39$, see Fig. 3.1.D), demonstrating that – as predicted under upweighting accounts – expected target events were perceived to be more forceful than unexpected events. This finding suggests that when tactile events can be explicitly predicted by action on the basis of statistical probabilities, perception of tactile events is biased *towards* expectations. This finding is therefore difficult to reconcile with claims in the action literature that predicted percepts should be perceptually downweighted, and is inconsistent with findings of tactile attenuation thought to arise from being able to predict the sensory consequences of action (e.g., Bays et al., 2005, 2006; Blakemore et al., 1998; Kilteni et al., 2019).

Experiment 2 therefore demonstrates for the first time that action predictions – formed on the basis of statistical learning – enhances perception for expected, relative to unexpected tactile events. Experiment 3 aimed to provide a replication of this study but adapted some aspects of the set-up to control for additional potential confounds in the paradigm – ensuring that effects are not specific to these action types or the result of predictions formed about the reference stimulus or the imperative cue.

3.2. Experiment 3

Experiment 3 was designed to provide a conceptual replication of the statistical learning paradigm of Experiment 2, while using a set-up more closely aligned with typical action paradigms, whereby one always makes a movement *towards* another effector (e.g., like in Experiment 1 and typical action paradigms). This induces a more ‘naturalistic’ mapping between action and sensory outcomes, which has been suggested to be important for observing tactile attenuation (Kilteni & Ehrsson, 2017b). Participants’

right hand was therefore vertically aligned with their left hand, such that movements of their right index/middle finger were directed towards their left index/middle finger, respectively (Fig. 3.1.B). This also ensured that the effects in Experiment 2 were not specific to those action types and determined whether similar effects could be observed with actions that are always made towards another effector. Similarly to in Experiment 3, one of two actions could predict one of two stimulation locations on the basis of conditional probabilities established in the training phase.

Two further changes were made in Experiment 3 to remove potential confounds. In Experiment 2, a cue instructed participants which action they should perform. In principle, the expectation effects observed in Experiment 2 may have resulted from cue-effect learning rather than action-effect learning. Experiment 3 therefore removed these cues and required free selection of action. The explicit reference stimulus was also removed in Experiment 3 to eliminate the possibility that expectation alters perception of the reference rather than target event. Information regarding the execution of actions was also recorded in this experiment to examine whether there are any differences in action durations or distance traversed across expected and unexpected events. It was hypothesised that if learned action-outcome associations influence tactile perception similarly regardless of action type, and when some potential confounds are controlled for, tactile events should be perceived more intensely than unexpected events – consistent with upweighting accounts and the results of Experiment 2.

3.2.1. Methods

Participants

Thirty new participants (22 female, mean age = 24.3 years [$SD = 4.34$]) were recruited like in Experiments 1 and 2. Three participants were replacements for those who could

not complete the perceptual discrimination, and a further six for those who were unable to follow instructions concerning movement performance (>20% movement errors on average e.g., where the appropriate action could not be detected by the motion tracker).

Procedure

The following changes were made relative to Experiment 2. Independent solenoids were attached to both left index and middle fingers via adhesive tape (diameter of metal rod = 4 mm; diameter of solenoid = 18 mm; TactAmp 4.2 Dancer Design), positioned so that the metal rod of the solenoid sat on the apex of the fingertip. Participants' hands, and therefore index and middle fingers, were aligned with each other (see Figure 3.1.B). Actions were freely selected and the frequency of index and middle finger movements was monitored to ensure approximately equal numbers of both action types. At the start of each trial, participants freely selected to make a movement with either their right index or middle finger. Participants' actions (e.g., right index downwards movement) were still perfectly predictive of the location of tactile events (e.g., left index finger) during training, and this contingency was again degraded to 66.7% in the following test session. After a 300 – 500 ms delay, participants were asked whether they perceived the test tap to be more or less forceful than the average tap. An example of the average force was presented to each finger once at the end of short breaks every 21 trials (NB: the average force was identical to the reference force intensity in Experiments 1 & 2). Responses were recorded by pressing a foot pedal positioned at either 45 (for stronger) or 90 (for weaker) degree angles relative to their right foot, to account for any spatial biases resulting from positioning foot pedals as 'left' and 'right'. The next trial started following 1000 ms, and an example of the average tap was presented to each finger once every 21 trials.

The experiment consisted of two training sessions followed by a test session, all now occurring on the same day. There were 210 trials in each session. During the training blocks, movements were perfectly predictive of tactile events, where half of the participants experienced a congruent action-stimulus mapping (e.g., move index finger, receive index stimulation) and the other half experienced an incongruent mapping (e.g., move index finger, receive middle finger stimulation). In the first training block participants responded to the question “Tap on index or middle finger?” with a yes/no response with the foot pedals. In the second training block they were asked about the force, similarly to in Experiment 2 and in subsequent test blocks.

3.2.2. Results and Discussion

Action execution via motion tracker

All conditions in this chapter relate to No Contact set-ups where actions were recorded using a motion tracker that measured and tracked the coordinates of fingers and hand positions throughout trials. Tactile stimulation was triggered once the task-relevant right-hand finger had traversed at least 20 mm from its initial starting position in the y-axis. In Experiment 3, these coordinates were stored and movement distance and duration were calculated relative to the starting coordinates. On average, participants’ actions traversed a distance of -25.11 mm ($SD = 1.81$) in durations of 627 ms ($SD = .19$). Distance traversed (Expected: $M = -25.20$ mm, $SD = 1.89$, Unexpected: $M = -25.02$ mm, $SD = 1.75$, $t(29) = -1.38$, $p = .178$, $d = -.25$) and durations (Expected: $M = 626$ ms, $SD = .19$, Unexpected: $M = 628$ ms, $SD = .20$, $t(29) = -.23$, $p = .817$, $d = -.04$) were similar in both conditions – an unsurprising similarity given that participants were unaware of the outcome at the time of action execution. It is likely that this data is also applicable to actions performed in Experiments 1 and 2, given that movements were observed by the experimenter and corrected via instruction across all experiments.

Expectation analyses

Like in Experiment 2, PSE values were lower for expected ($M = 4.08$, $SD = 1.00$) than unexpected ($M = 4.28$, $SD = .99$) events ($t(29) = -2.56$, $p = .016$, $d = .47$, see Fig. 3.1.D). This result reveals that expected tactile events were perceived to be more forceful than unexpected events, consistent with the results of Experiment 2 and upweighting theories of perception.

The findings of Experiments 2 and 3 suggest that tactile events that can be learned to be expected on the basis of action, are perceived more forcefully than unexpected events, and do not depend on the specific action type or do not reflect predictions formed about the reference or imperative cue. Furthermore, they suggest that arbitrary action-outcome mappings, even if considered ‘unnatural’ (e.g., upwards right finger motion that predicted force to left hand below in Experiment 2), do not seem to influence perception differently from actions that are always made towards other effectors (Experiment 3) – likely reflecting the statistical learning that took place between actions and outcomes regardless of action type. These findings would be inconsistent with assumptions of action researchers that predictive influences of action depend on ‘naturalistic’ self-touch mappings (Kilteni & Ehrsson, 2017b). Nevertheless, these results suggest that action has the capability of influencing perception of predictable tactile events in comparable ways to reports of predictive influences outside of action contexts and in line with upweighting theories.

3.3. Computational modelling

The findings presented in this chapter suggest that expected action effects are perceived with greater intensity than unexpected action effects, consistent with predictive upweighting theories of perception. These theories propose that it is adaptive for observers to combine sampled sensory evidence with prior knowledge – biasing perception towards what we expect (Yuille & Kersten, 2006). This may be achieved mechanistically by altering the weights on sensory channels, increasing the gain of expected relative to unexpected signals (de Lange et al., 2018; Summerfield & de Lange, 2014), and increasing the likelihood that expected signals are detected and perceived more intensely than unexpected signals. However, an alternative explanation for the findings is that expectation effects are a product of biasing in response-generation circuits, such that action biases people to *respond* that events are more intense when they are expected, rather than biasing *perception* itself (see De Lange, Rahnev, Donner, & Lau, 2013; Firestone & Scholl, 2016).

These different kinds of bias can be dissociated in computational models that conceptualise perceptual decisions as a process of evidence accumulation. Perceptual biases are thought to grow across time – every time response units sample from perceptual units they will be sampling from a biased representation, therefore increasing the magnitude of biasing effects across a larger number of samples (Urai et al., 2019; Yon et al., 2020). In contrast, response biases are thought to operate regardless of current incoming evidence and be present from the outset of a trial (Leite & Ratcliff, 2011). According to this logic, the decision process can be modelled using drift diffusion modelling (DDM, Ratcliff & McKoon, 2008) to identify the type of biasing that is generating effects. DDMs conceptualise two-choice decisions (e.g., ‘stronger or weaker than average?’) as a noisy process of sequentially sampling sensory evidence to

compute a decision variable (Ratcliff & McKoon, 2008). Varying one parameter, z , of the DDM allows us to establish whether action expectations shift the starting point of evidence accumulation towards a response boundary (Leite & Ratcliff, 2011), therefore reaching one boundary ('stronger') faster than the other ('weaker'; 'start biasing', see Fig. 3.2.A). On the other hand, varying the 'drift biasing' parameter (db ; Fig. 3.2.B) allows us to establish whether the rate of accumulation is faster towards one boundary than the other, as may be expected if the sensory units are biased (Mulder et al., 2012).

3.3.1. Methods

Hierarchical DDMs (hDDMs) were fit to participant choice and reaction time (RT) data from Experiment 3 using the hDDM package implemented in Python (Wiecki, Sofer, & Frank, 2013; it was not possible to do the same for Experiment 2 as RTs were not collected). In the hDDM, model parameters for each participant are treated as random effects drawn from group-level distributions, and Bayesian Markov Chain Monte Carlo (MCMC) sampling is used to estimate group and participant level parameters simultaneously. Four different models were specified: 1) a *null* model where no parameters were permitted to vary between expected and unexpected trials; 2) a *start bias* model where the start point of evidence accumulation (z) could vary on expected and unexpected trials; 3) a *drift bias* model where a constant added to evidence accumulation (db) could vary between expected and unexpected trials; 4) a *start + drift bias* model where both parameters could vary between expected and unexpected conditions. All models were estimated with MCMC sampling, and parameter estimates were converged after 30,000 samples ('burn-in' = 7,500). Models were compared using deviance information criteria (DIC) as an approximation of Bayesian model evidence, a common method used to determine model fit. Lower DIC values relative to a baseline, or null, model are indicative of a better model fit.

As well as examining the model-fit statistics, a posterior predictive check was conducted using the hDDM package to establish how well each model was able to reproduce the patterns in the data. The posterior model parameters for the *start bias*, *drift bias*, and *start + drift bias* models were used to simulate a distribution of 500 reaction times and choices for each trial for each participant under each model. From this simulated data, the probability that a ‘stronger than average’ response was given was calculated at each intensity level, separately for expected and unexpected trials. This allowed simulated psychometric functions to be modelled for expected and unexpected trials, exactly as was done for empirical decisions. Performing this procedure for each model yielded separate simulated expectation effects (Unexpected PSE – Expected PSE) for each participant under the *start bias*, *drift bias* and *start + drift bias* models.

3.3.2. Results and Discussion

Fitting the DDM to the behavioural data found that the model including both start and drift biases provided a better fit (DIC relative to null = -234.8) than both the start bias (DIC relative to null = -191.06) and drift bias (DIC relative to null = -8.62) models. This finding may suggest that observed biases are a product of both start and drift rate biasing. However, although the DIC measure does include a penalty for model complexity, it is thought to be biased towards models with higher complexity (Wiecki et al., 2013) and it indeed favoured the most complex model here.

To examine whether any of the PSE expectation effect is generated by sensory biasing – rather than possible additional contributions of response biasing – a posterior predictive check was conducted to evaluate how well simulated data from each of the models could reproduce key patterns in the data. Correlations were calculated to quantify how

well simulated expectation effects reproduced the observed empirical expectation effects, which revealed significant relationships for all three models (Start bias model: $r_{30} = .39, p = .034$; Drift bias model: $r_{30} = .43, p = .017$; Start + Drift bias model: $r_{30} = .53, p = .003$, see Fig. 3.2.C).

In order to establish whether drift biasing accounted for any further variance in expectation effects than start biasing alone, a stepwise linear regression was conducted to predict the empirical expectation effect (Unexpected PSE – Expected PSE). In the first step, the simulated PSEs from the start bias model were included to predict the empirical PSE effect. The simulated start bias data was able to predict the empirical PSE expectation effect ($R^2 = .15, F(1,28) = 4.96, p = .034$). In the second step, the simulated PSEs from the drift bias model were included as an additional predictor of the empirical PSE effect. The regression model remained significant with this addition ($R^2 = .32, F(2,27) = 6.34, p = .006$), and importantly provided a significant improvement to the model fit ($Fchange(1,27) = 6.72, p = .015$). This analysis reveals that a model implementing a drift biasing mechanism better predicts empirical effects of expectation on perceptual decisions, by explaining unique variance in participant decisions that cannot be explained by response biasing. Drift biasing reflects how sensory evidence is accumulated and is what would be anticipated if the sensory representations were themselves biased – as hypothesised under the Bayesian account.

However, it is also important to note that there has been a recent suggestion that pre-activation mechanisms – shown to enhance expected neural sensory representations ahead of predictable input (e.g., Blom et al., 2020; Kok et al., 2017) – could also be modelled by starting point biases (Feuerriegel et al., 2021). Feuerriegel and colleagues (2021) reason that shifts in starting points of evidence accumulation towards the

expected boundary could be reflective of an earlier onset of evidence accumulation, where decision processes are biased by sampling from these pre-activated sensory representations – and manifest in faster responses to expected than unexpected stimuli (Feuerriegel et al., 2021). This reasoning may be especially applicable to the present results because the interval between stimulus presentation and the response window meant that sensory evidence could be accumulated up to 500 ms before the onset of decisional processes. On the other hand, response biases present from the outset of the decision process could come about by a decisional shift that is independent from the specifics of the sensory evidence. It is perhaps more difficult to imagine how this would be implemented, considering that the particular response question (of force intensity) was independent of expected stimulus features (force location) and would thereby require a decision shift towards responding ‘stronger’ on expected trials regardless of intensity. However, further tests would be required to definitively draw conclusions regarding the nature of start biasing observed in this experiment. Nevertheless, these findings demonstrate evidence that perceptual biases likely underlie Experiment 3’s expectation effects, suggesting that the sensory representations were themselves biased.

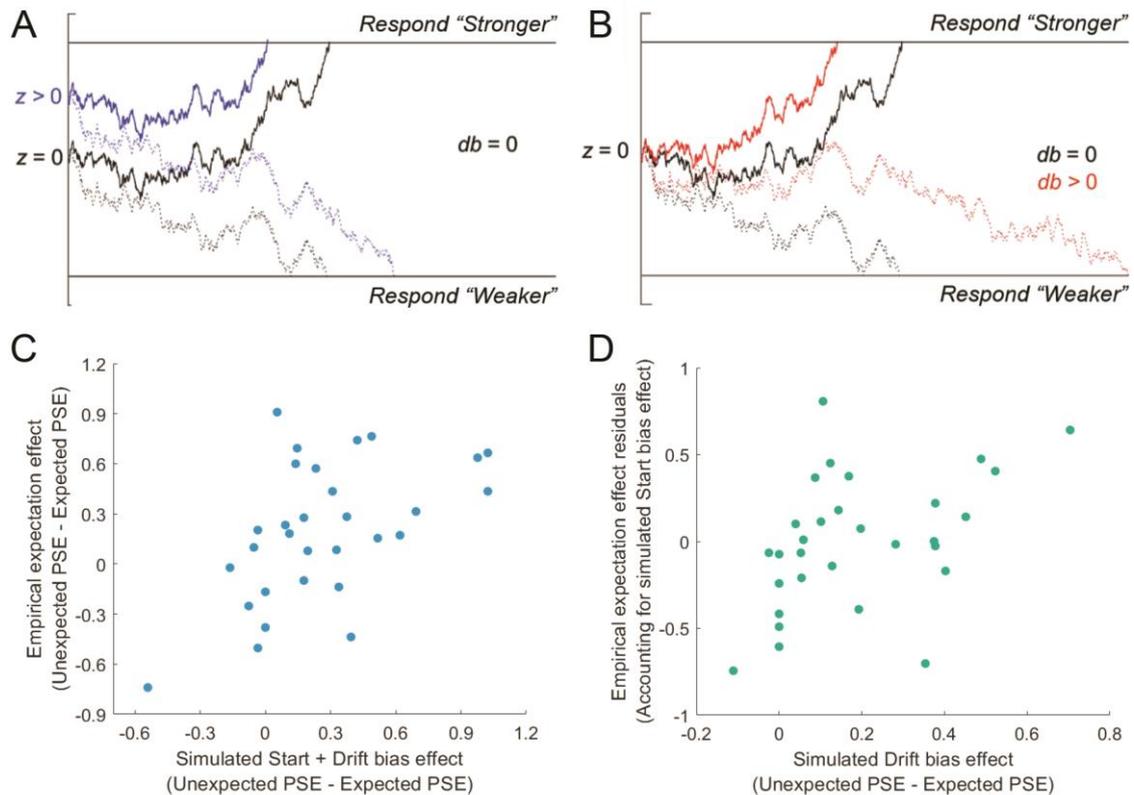


Figure 3.2. (A) In an unbiased decision process (black lines), sensory evidence is accumulated towards the upper response boundary when stimuli are stronger than average (solid lines) and towards the lower response boundary when weaker than average (dotted lines). The starting point of evidence accumulation could however be shifted by decision circuits towards the upper boundary for expected events from the outset of the response window (influencing the parameter z ; blue lines - Start bias model). (B) Alternatively, altering the weights on sensory channels could bias evidence accumulation in line with expectations (influencing parameter db ; red lines – Drift bias model). Adapted from Yon et al., (2020). (C) Simulated Start + Drift bias (winning DIC model) expectation effect plotted against the empirical expectation effect, showing a significant positive correlation. (D) Simulated Drift bias expectation effect plotted against the empirical expectation effect accounting for simulated Start bias effects (plotted as the residuals from a model where the simulated Start bias effect predicts the empirical effect), again showing a significant positive correlation. Importantly, the regression analysis revealed that drift biases accounted for significant additional variance once accounting for start biases. All expectation effects were calculated by subtracting Expected PSEs from the Unexpected PSEs.

3.4. Experiment 4

Thus far, Experiments 2 and 3 have found evidence in support of predictive upweighting theories – a process that is thought to aid veridicality of perceptual experiences – and reveal that expected tactile action-outcomes are perceived with greater intensity than unexpected outcomes. These results are however harder to reconcile with current accounts in action domains that propose that predicted action outcomes are downweighted, serving the functional role of optimising informativeness for updating current models and guiding corrective action.

One might interpret these data as reflecting that upweighting mechanisms likely operate in a domain general fashion and that downweighting theories should be discarded.

However, such a conclusion would be premature on the basis of the evidence. Firstly, although some downweighting effects could have alternative explanations, other studies provide more convincing evidence that unexpected action outcomes are perceived more intensely or reliably than expected events (Desantis et al., 2014; Roussel et al., 2014; Yon & Press, 2017). Secondly, we need our experiences to be informative as well as veridical and therefore need to consider how the brain might achieve this functional aim (Press et al., 2020b). A potential resolution to the debate comes from a recently proposed ‘opposing process’ theory. Under this account, perception is initially biased towards what we expect in order to generate experiences rapidly that are more veridical, but if events are presented which generate particularly high surprise, later processes adapted to subserve learning highlight these events, outweighing processing of expected events (Press et al., 2020b). The earlier process is thought to operate the moment predictions are formed and therefore acts to pre-emptively activate expected sensory representations, while the later process will only highlight the most informative unexpected events and would therefore only occur for a subset of surprising events.

Therefore, this theory proposes that both upweighting and downweighting can occur with differing functional roles and across different timescales. Although largely yet to be tested, some preliminary data in visual perception supports this idea, and reveals that expected action-outcomes are perceived as brighter at early time points (50 ms after stimulus presentation), yet at later probes (200 ms later) this bias flips over such that unexpected outcomes are now perceived more intensely than expected events (Yon & Press, 2017; see also Press et al., 2009).

One possible reconciliation for the current debate is therefore that both functional processes occur but at different timescales, where methodologies adopted in support of either theory are often only optimal to detect either one of these processes. For example, Experiments 2 and 3 are consistent with the operation of the earlier upweighting process given that tactile events were punctate (30 ms), however, it is plausible that hypothetical later surprise-based processes could not have influenced perception because these processes would operate after stimulus offset. Experiment 4 therefore aims to address the claim that predictive opposing processes operate at different time points in perceptual experience, to examine whether such a hypothesis could help explain the discrepancy across literatures examining predictive influences on perception. Tactile events were thereby presented at either short (comparable to events in Experiments 2 and 3) or long durations, and again, could be either expected or unexpected on the basis of learned action-outcome associations.

The set-up of Experiment 4 was almost identical to the congruent mapping group of Experiment 3 demonstrating strong expectation effects, but now presents different duration tactile vibrations. Delivering vibrations allows manipulation of tactile events longer than 50 ms – which was not previously possible when presenting solenoid forces

via a single square-wave – in order to test whether predictive effects occur at different timescales. Under the opposing processes theory, early upweighting would be hypothesised to enhance perception of expected events relative to unexpected events at short durations – similarly to as observed in Experiments 2 and 3. However, the theory would hypothesise that later downweighting processes might operate to highlight perception of unexpected relative to expected events when longer durations events are presented. Therefore, the opposing process theory would expect, at minimum, that any enhancement effect observed for short events should be reduced for longer events, but in principle such effects may disappear or even reverse.

3.4.1. Methods

Participants

Thirty new participants (16 female, mean age = 26.80 years [$SD = 6.94$]) were recruited from Birkbeck, University of London and were paid a small honorarium for their participation. Five participants were replacements for those who were unable to follow instructions concerning movement performance (>20% movement errors on average) or who experienced technical issues.

Procedure

Additional changes in Experiment 4 relative to Experiment 3 are described here. The hands were positioned using the same set-up as Experiment 3. The target stimulus was now presented as a vibration (25 Hz) for either a short (50 ms) or long (500 ms) duration. Stimulus intensity for short vibrations were identical to those used in Experiments 1-3, and these intensity levels were matched during piloting to perceived intensities at long vibrations. This piloting also enabled generation of target stimulus

intensities that were perceived as equal as possible in precision across the two duration sets.

The experiment consisted of two training sessions and a test session occurring on the same day. There were 42 trials for each training session, and 420 trials in the main test session. Participants only ever received a congruent mapping between their action (e.g., right index) and tactile outcome (e.g., left index) at training. Blocks alternated between short and long duration events and this order was counterbalanced across all participants. Implicit reference events were presented at short or long durations at the start of each mini-block, for the respective duration.

3.4.2. Results and Discussion

PSE values for the factors Expectation (expected and unexpected) and Duration (short or long vibrations) were analysed in a 2 x 2 mixed ANOVA. The PSE analysis revealed no main effect of Expectation, $F(1, 29) = 3.51, p = .071, \eta_p^2 = .11$, but a main effect of Duration, $F(1, 29) = 8.17, p = .008, \eta_p^2 = .22$, and most interestingly a significant interaction between Expectation and Duration, $F(1, 29) = 13.00, p = .001, \eta_p^2 = .31$. This interaction was driven by lower PSEs for expected trials ($M = 3.61, SD = .92$) than unexpected trials ($M = 4.03, SD = .94$) when vibration durations were short ($t(29) = -3.37, p = .002, d = .62$), but no difference between expected ($M = 3.33, SD = .89$) and unexpected ($M = 3.33, SD = .91$) trials when vibrations were long ($t(29) = -.01, p = .993, d = <.01$; see Fig 3.3.A).

These results revealed that short tactile events expected on the basis of action are perceived as more forceful than unexpected stimuli, consistent with the results of Experiment 2, 3 and predictive upweighting theories. Interestingly, this effect was not present when longer tactile events were presented, as there were no differences between

expected and unexpected conditions. The fact this effect disappears with longer duration stimuli is consistent with assumptions of the opposing process account because it would propose that later processes increase the perceived intensity of unexpected events (Press et al., 2020b).

However, there are some possible alternative explanations. The resulting null effect at long durations may reflect perceptual ceiling effects, considering that longer tactile events provide more sensory evidence. It is important to note that precision of functions for both durations was matched during piloting and showed no difference across conditions in Experiment 4, therefore this conclusion is perhaps unlikely. However, participants tended to report target long duration events as more forceful than short events relative to the respective reference events, which likely reflects greater sensory evidence of a force even if the precise intensity could not be determined any more easily in this condition.

An explanation consistent with assumptions of the opposing process theory could be that long duration null effects reflected the operation of both early (enhancing) and later (attenuating) processes within this perceptual time window. Considering that these mechanisms are opposing in nature, it is likely that a later process highlighting unexpected events would cancel out any early highlighting of expected outcomes when events are presented at durations that would encompass both processes – thereby resulting in a flat expectation effect overall. However, because the temporal window for such a flip has not been tested in tactile perception, this explanation is only speculative and clarification should be provided in future studies by presenting tactile events at different delays, perhaps closer to those studies in vision (e.g., 200 ms; Yon & Press, 2017). Effects observed at different time windows may provide insight concerning the

temporal evolution of potentially distinct processes and help to establish whether there are indeed opposing predictive processes influencing perception over time.

In conclusion, Experiment 4 importantly replicates findings from Experiments 2 and 3 that short expected action outcomes are rated as more forceful than short unexpected forces. It also demonstrates that these perceptual advantages are abolished when longer duration events are presented, and may indicate that alternative perceptual mechanisms operate at later points in time.

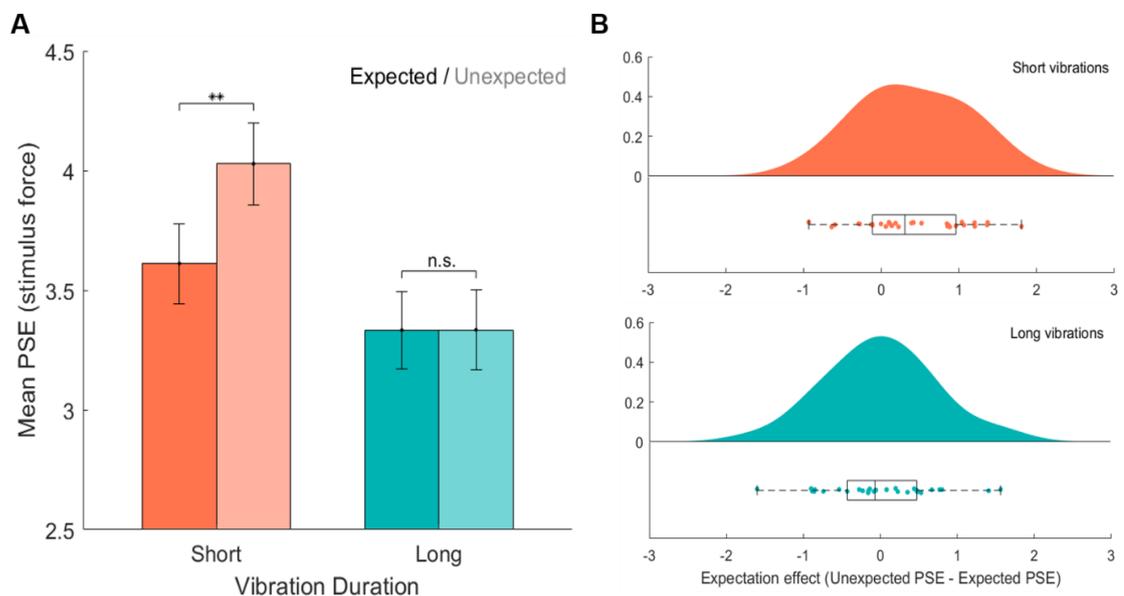


Fig 3.3. (A) Mean PSEs (\pm SEM) were lower for expected than unexpected trials for short vibration durations but not different for long duration vibrations (** $p < .01$). Lower PSEs indicate more intensely perceived tactile events. (B) PSE expectation effects (Unexpected – Expected) plotted with raincloud plots (Allen et al., 2019) displaying probability density estimates (upper) and box and scatter plots (lower), for short (top) and long (bottom) duration conditions. Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive expectation effect values indicate more intensely perceived expected events relative to unexpected events.

3.5. General Discussion

This chapter examined whether sensorimotor prediction attenuates perception of tactile events, as is widely assumed in the action literature (Blakemore et al., 1998; Dogge, Custers, & Aarts, 2019; Fiehler et al., 2019; Kilteni & Ehrsson, 2017a), or instead whether predicted events may be perceptually upweighted – in line with theories in the wider sensory cognition literature (de Lange et al., 2018; Kersten et al., 2004). A version of the force judgement paradigm typically adopted in the action literature was combined with typical manipulations of prediction according to statistical probabilities from the broader sensory cognition literature (e.g., Cheadle et al., 2015; Han et al., 2019; Kok et al., 2012, 2013, 2014; Kok & Turk-Browne, 2018; Richter & de Lange, 2019). Experiments 2-4 examined the relative intensity of tactile events expected or unexpected on the basis of actions, thereby rendering a manipulation more equivalent to those found outside of action. All three experiments found that expected tactile action outcomes were perceived *more*, not less, intensely than those that were unexpected, in particular when events were short in duration (Experiment 4). Computational modelling suggested that expectations alter the way sensory evidence is integrated in Experiment 3 – increasing the gain afforded to expected tactile signals.

These findings are consistent with upweighting perceptual accounts from outside of action domains, proposing that we use prior knowledge to bias our percepts towards those that are more probable (de Lange et al., 2018; Kersten et al., 2004; Yuille & Kersten, 2006). Such accounts propose that expectations that are learned through the statistical regularities of our environment are combined with sensory evidence to generate veridical perceptual interpretations of our noisy sensory world. These theories have been developed outside of action research, but in fact have already been demonstrated to be consistent with how expectation influences *visual* processing during

action (Christensen et al., 2011; Dogge, Custers, Gayet, et al., 2019; McDonough et al., 2019; Yon et al., 2018, 2020; Yon & Press, 2017). In touch, some evidence has suggested that movement preparation can improve perceptual identification accuracy of tactile braille stimuli (van Ede et al., 2015), but there is limited evidence to suggest that *predictive* influences of action on tactile perception operates in a similar manner to in visual perception. The present findings therefore indicate for the first time that predictive upweighting mechanisms operate similarly in touch, suggesting that sensorimotor prediction does not influence touch perception in a special fashion relative to other sensory modalities.

The findings are harder to reconcile with predominant downweighting theories within the action literature (Blakemore et al., 1998; Dogge, Custers, & Aarts, 2019; Kilteni & Ehrsson, 2017a), that propose expected action outcomes are attenuated. It is therefore essential to consider how the present findings fit amongst the multitude of data cited in support of tactile attenuation. A large body of work has attributed attenuation to predictive processes in humans as well as in a variety of other species and sensory systems. For example, attenuating internally-generated electric fields in Mormyrid fish improves detection of prey-like stimuli, a finding thought to be a consequence of predicting self-generated sensory input (Enikolopov et al., 2018). Similarly, virtual reality trained mice show suppressed auditory responses to self-produced tones generated by treadmill running (Schneider et al., 2018) or licking behaviours (Singla et al., 2017), compared to no movement. As discussed, in humans, studies measuring the perceived force of a tactile stimulus during movement show similar attenuation effects during movement relative to no movement, and commonly attribute these effects to predictive mechanisms (Bays et al., 2005, 2006). However, such studies demonstrating

tactile attenuation have rarely demonstrated whether underlying mechanisms operate according to stimulus probabilities (de Lange et al., 2018).

There are a number of non-predictive mechanisms which could mask predictive influences on perception when such statistical manipulations are not performed, and based upon current and previously discussed findings it is possible that some effects are generated by sensory gating mechanisms, yet others might alternatively shape perception according to event repetition – given that repetition is frequently confounded with expectation (Feuerriegel et al., 2020). In fact, a recent study by Kilteni et al. (2019) aimed to determine whether the attenuating influence of action on tactile perception indeed reflected the operation of predictive mechanisms according to statistical probabilities. In an extension of the force judgement paradigm, the authors trained participants that sensory events would be delayed by 100 ms, and demonstrated an inverted typical attenuation effect such that delayed (but expected) events were now perceived less intensely than simultaneous (but unexpected) events. They therefore concluded that we do indeed downweight perception of tactile effects predicted by action – as has been assumed for the last few decades but never explicitly demonstrated (Kilteni et al., 2019). However, while the environmental statistics were altered in Kilteni et al.'s (2019) study to make delayed events more probable, this training protocol likely confounds prediction with repetition. Participants in the test group were trained in only one repeated mapping – consisting of a 100 ms action-outcome delay – and are subsequently presented with events at that delay or at no delay. Given that repeated presentation of stimulus features can attenuate perception of those features (Ofen et al., 2007) and neural sensory processing (Grill-Spector et al., 2006), downweighting effects observed in this study may be determined by repetition of the

100 ms delay rather than prediction of it, especially given that every response trial was preceded by an extra five repeated training trials.

The distinction between mechanisms underlying repetition and prediction effects has been scrutinised widely for the last decade and a number of differences emerge. For instance, some studies have demonstrated that expectation effects (e.g., visual mismatch response) are only observed when the presented stimulus is not a repeat of the previous stimulus, and therefore manipulations of repetition appeared to reduce orthogonally manipulated effects of expectation (Feuerriegel et al., 2018, 2019). However, other studies have shown that when repetition and prediction are orthogonalised by sometimes rendering stimulus alternations predicted, repetition and prediction neural effects exhibit distinct time-courses (Todorovic & De Lange, 2012). More recently, evidence has emerged that repetition neural effects are likely mediated by mechanisms functionally distinct from prediction effects – with debates surrounding whether prediction effects emerge via neural sharpening (Kok, Jehee, et al., 2012; Yon et al., 2018) and repetition effects via local neural scaling (Alink et al., 2018; but see Alink et al., 2020; Ramírez & Merriam, 2020).

Importantly however, the present data should not be taken to reflect that predictive attenuation cannot occur (especially given the importance of generating perceptual experiences that are informative across domains; see Press et al., 2020b), but rather that the mechanisms generating attenuation are unlikely to operate as assumed across the last few decades. In other words, the assumed mechanisms underlying predictive tactile attenuation require pre-emptive downweighted perception of expected events – perhaps due to subtraction of the prediction from actual input (Wolpert & Flanagan, 2001) – but such a mechanism is hard to reconcile with the present findings demonstrating

upweighted sensory gain of predicted action outcomes. These data more likely suggest that purported predictive mechanisms must have the capability of generating upweighting as well as downweighting, but under different circumstances. As outlined in section 3.4., one possible resolution to the current debate explains how this may be achieved via opposing processes with differing roles. The enhancement effects found in Experiments 2-4 are consistent with this proposition given that tactile events were punctate (30 ms), and so hypothetical later surprise-based processes may not have had sufficient time to influence perception because they would operate after stimulus offset.

The results of Experiment 4 provide some evidence in support of a later attenuating process because expectation enhancements were eliminated for long duration events. For instance, it is possible that this null effect was produced by both early and late perceptual processes – highlighting expected in the former and unexpected events in the latter – operating within this time window and therefore cancelling each other out. However, it is hard to draw strong conclusions from a null effect in this condition, and therefore we cannot be sure that it reflected the operation of opposing processes. Future work is required to test assumptions from the opposing processes theory that upweighting and downweighting mechanisms can occur in different contexts and at different timescales across perceptual experience. Importantly, such work should provide further reconciliation to the prediction paradox by determining whether the opposing process theory could account for the discrepancies across findings in the literature.

Nevertheless, the important point to note from this set of tactile experiments is that sensorimotor prediction can increase the perceived intensity of tactile events similarly to studies examining visual perception both inside and outside of action domains, and

consistent with upweighting accounts of perception. These findings suggest that action predictions formed on the basis of statistical regularities influence tactile perception via qualitatively similar mechanisms to prediction mechanisms reported in the broader sensory cognition literature, and seemingly regardless of the sensory domain or the nature of the predictive cue.

Chapter 4: The nature of neural prediction mechanisms in action

So far it has been demonstrated that predictions formed on the basis of action can bias tactile perception towards expected events, increasing their perceived intensity relative to unexpected events (Experiments 2-4). This is consistent with some findings in visual perception revealing that actions elicit more intense perception of congruent visual outcomes (Yon & Press, 2017, 2018) and increase detection of such events (Yon et al., 2020) compared to incongruent outcomes. However, comparability between the reported tactile experiments and these visual studies is unclear because prediction was not manipulated according to statistical regularities in the latter. Congruency manipulations of visual action outcomes are assumed to be predictive due to the fact that congruent outcomes are statistically more likely to follow actions than incongruent outcomes (Cook et al., 2014). However, it is undetermined whether congruency effects are indeed reflective of the operation of predictive mechanisms, and therefore it is uncertain whether these visual effects reflect predictive processes.

The prediction paradox also extends to neural literatures aiming to explain underlying sensory mechanisms – with the action literature primarily reporting support for cancellation accounts in visual and somatosensory cortices, while the broader sensory cognition literature tends to report support for upweighting accounts in early visual regions. Across all domains, activation is typically lower in sensory brain regions for predictable events relative to less predictable events. Studies in action domains conclude that these effects provide evidence for predictive downweighting mechanisms, because activity associated with sensory consequences of action is suppressed to allow processing of unexpected action outcomes (Blakemore et al., 2001; Kilteni & Ehrsson, 2020a; Shergill et al., 2013; Stanley & Miall, 2007; Wolpe et al., 2018). Implemented

neurally, cancellation accounts propose that an ‘efference copy’ concerning the sensory consequences of action is passed to the appropriate sensory brain regions (Wolpert & Flanagan, 2001), resulting in a suppression of neural activity in expected sensory units and allowing us to become especially sensitive to informative unexpected action outcomes. Specifically, this suppression is thought to be largest in sensory populations tuned towards the expected outcome, and would predict a relative decrease in the contrast of activity patterns evoked by expected stimuli than unexpected stimuli (de Lange et al., 2018; Press et al., 2020b; see Fig 4.1.A and 4.1.C).

However, it has been proposed outside of action that expectation suppression could instead reflect the operation of ‘sharpening’ predictive mechanisms. Kok et al. (2012) demonstrated with MVPA that pattern classifiers trained to discriminate between stimulus orientations on the basis of V1 activation performed with superior accuracy for expected than unexpected events – revealing improved stimulus-specific representations of expected signals (Kok, Jehee, et al., 2012). Sharpening theories have gained increasing popularity in the last decade, and hypothesise that predictions evoke a ‘pre-activation’ of stimulus-specific representations of expected events while activity in sensory populations tuned towards unpredicted events is suppressed – perhaps via lateral inhibition. This process results in a sharpened representation of expected outcomes (de Lange et al., 2018; Kok, Jehee, et al., 2012; Press & Yon, 2019; Summerfield & de Lange, 2014; Yon et al., 2018). Under this account and in contrast with the cancellation account, expectation suppression will be demonstrated in sensory populations that are least selective to the expected stimulus (e.g., populations tuned to horizontally oriented lines when vertical orientations are expected), and this process predicts an increase in the contrast of expected relative to unexpected activity patterns, to reflect ‘sharpened’ representations (de Lange et al., 2018; see Fig 4.1.B and 4.1.C).

Evidence for exactly how this sharpening mechanism might operate has primarily been observed outside of action contexts in studies that manipulate the statistical probability of expected and unexpected events – yet such manipulations are not comparable to those within action and so this difference might also account for conflicting findings in the neural literature. For example, a variety of studies examining prediction mechanisms in visual perception manipulate the probability that sensory cues predict visual events. Such studies reveal that expected events can be decoded from visual brain activity before they are presented (Kok et al., 2017) and can induce stimulus-specific predictive templates in the absence of bottom-up input (Aitken et al., 2020; Ekman et al., 2017; Kok et al., 2014). Additionally, prediction related sharpening in the visual cortex has been found to correlate with activity related to the predicted cue in the hippocampus – known to be involved in sequence learning – when there is a statistical relationship between the cue and outcome (Kok & Turk-Browne, 2018). These findings support the idea that predictions formed by statistical learning likely elicit pre-activated representations of expected outcomes in line with predictive upweighting accounts.

However, it is yet to be demonstrated whether the same neural mechanisms underpin perception of events predicted by action. Therefore, the experiments in this chapter examine whether predictive mechanisms in early visual regions operate in a similar fashion when sensory events are predicted via statistical association with action, or whether sensorimotor prediction mechanisms operate in a qualitatively distinct fashion during action. Experiment 5 initially examines the influence of sensorimotor prediction on visual perception in a statistical learning paradigm, which is then adapted in Experiment 6 to investigate the nature of neural predictive mechanisms during action.

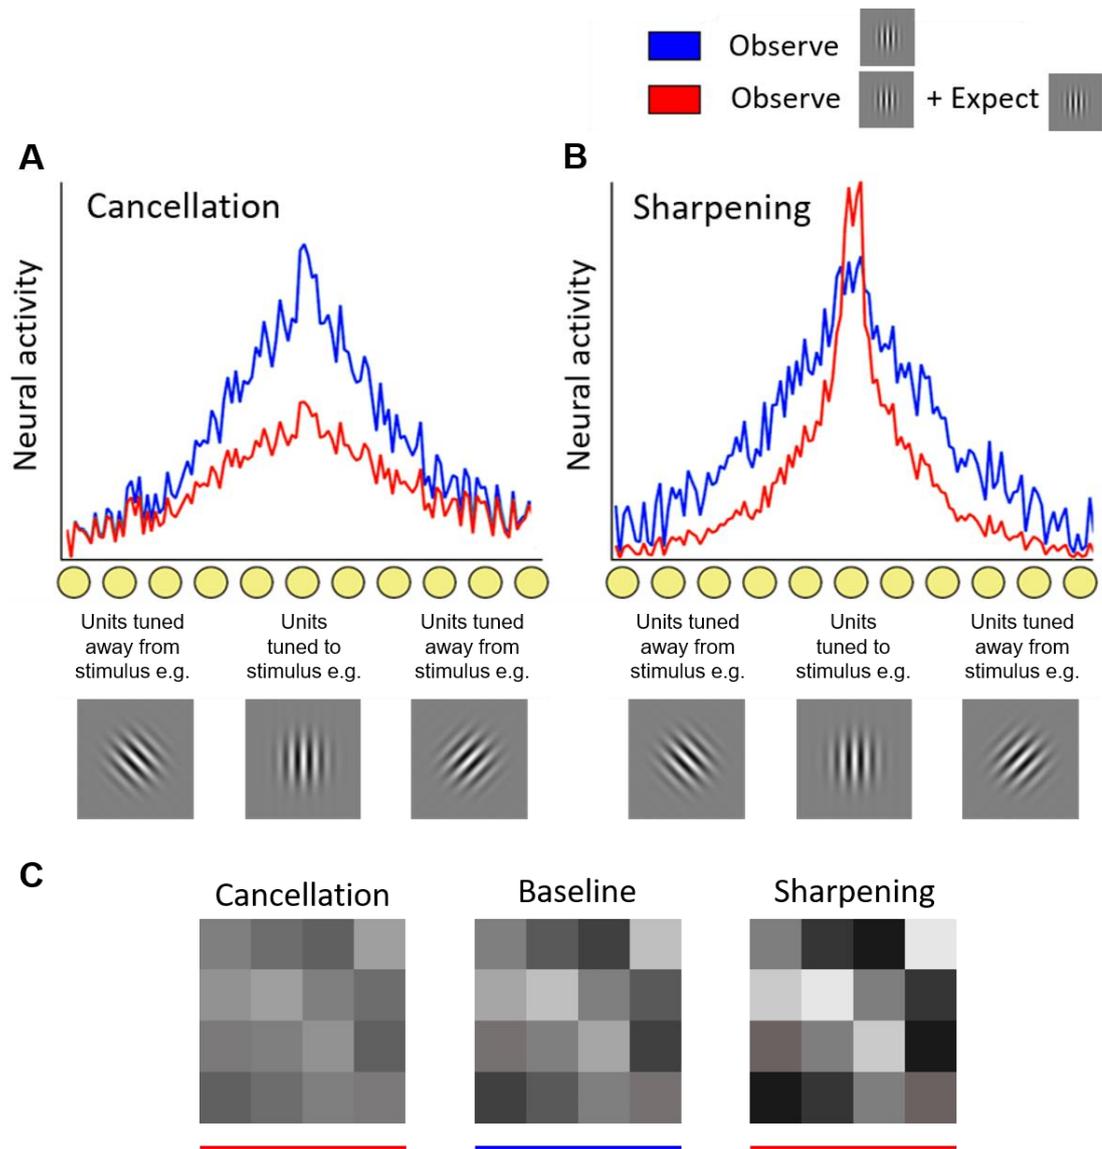


Figure 4.1. (A) Cancellation theories propose that expected action-outcome representations are suppressed in order to highlight processing of unexpected events. This expectation suppression will be largest for units tuned towards the expected stimulus. (B) Sharpening theories propose that we activate expected representations and relatively suppress processing in units representing other features or events. (C) These processes will result in reduced contrast in activity patterns for expected outcomes relative to unexpected outcomes in the former (cancellation), yet increased contrast in expected activity patterns in the latter to reflect sharpened sensory representations (e.g., high signal-to-noise ratio; adapted from Press et al., 2020).

4.1. Experiment 5

As discussed, some preliminary evidence in visual perception has suggested that action predictions bias perceptual decisions towards expected events, e.g., congruent action outcomes are perceived as more intensely (Yon & Press, 2017, 2018) and their presence is reported more frequently than unexpected outcomes (Yon et al., 2020). These findings are in line with predictive upweighting theories, similarly to the predicted touch experiments presented in this thesis. However, the predictive nature of these visual effects during action is less certain because the studies examine congruency. It is assumed that congruent action outcomes (e.g., index finger movements when moving one's index finger) can be expected on the basis of a lifetime of experience, but without manipulating statistical relationships within an experimental setting we cannot be sure of the nature of supporting mechanisms.

Experiment 5 therefore aimed to determine whether action predictions influence visual perception via similar mechanisms when they could be formed on the basis of statistical learning processes, similarly to those established outside of action. This experiment therefore trained participants on the association between an index or little finger action and the presentation of an oriented visual grating stimulus (clockwise or counter-clockwise), and degraded the contingency of this association at a test session 24 hrs later to measure responses to expected and unexpected events. This paradigm was intended to be utilised in Experiment 6 to determine the neural mechanisms of action predictions in an fMRI scanning procedure, and therefore Experiment 5 was predominantly concerned with establishing a behavioural effect that could provide evidence for action-outcome learning. Initially, evidence for learning through RT measures was assessed for expected and unexpected orientation judgements.

However, I was also curious as to whether evidence for a perceptual bias could be found, similarly to that has been demonstrated in paradigms outside of action. Namely, visual feature expectation has been manipulated to determine how perceptual sensitivity may be influenced by expectation in a signal detection paradigm (Cheadle et al., 2015). The authors manipulated feature prediction using a cue that indicated the probability of visual grating stimuli occurring on screen (e.g., a pink cue predicted a counter-clockwise tilted stimulus with 67% likelihood and a clockwise stimulus with 33% likelihood). They adopted a reverse correlation approach to quantify the extent to which increases in stimulus energy (e.g., for a clockwise stimulus) enhance the likelihood of a correct response (e.g., respond “clockwise”). The results interestingly revealed that decisional templates for expected features were shifted, such that participants’ responses became more sensitive to exaggerated versions of expected features e.g., to more extremely oriented clockwise features, suggesting that expectations bias representations of orientation towards extreme examples of expected categories.

This phenomenon is also known as a ‘peak shift’ effect, and was first discovered in animal learning studies (Wills & Mackintosh, 1998). For example, classic peak shift effects are apparent in set-ups that train animals to discriminate between two categories of stimuli, reinforcing responses to one and not the other. Such studies observe that conditioned responses become strongest to a stimulus that is neither of the trained stimuli but a more extreme version of the conditioned stimulus, that is, shifting the peak of the response gradient away from an undesirable stimulus and towards extreme versions of the desirable stimulus (Livesey & McLaren, 2017; Mackintosh & Little, 1969; Wills & Mackintosh, 1998). This phenomenon has also been observed in humans, for example in a study that required participants to judge the direction of a display of clockwise or counter-clockwise random-dot motion (relative to a reference

discrimination boundary line in between the two stimulus categories), reproduced motion directions are reportedly biased away from the discrimination boundary (Jazayeri & Movshon, 2007). One suggestion for the underlying mechanism of these effects is that each stimulus will activate a number of perceptual representations of different motion directions, and during perceptual decision-making representations tuned to the most remote or extreme representations (e.g., furthest away from the discrimination boundary) are given the highest weight yet representations tuned to orientations closest to discrimination boundaries are given the least weight (Jazayeri & Movshon, 2006). This is likely because those close to discrimination boundaries are activated equivalently for both events being distinguished. Under such accounts and in line with the findings of Cheadle et al. (2015), if a 45° oriented stimulus is expected while a 135° stimulus is unexpected, representations of more extremely tilted (e.g., $\sim 40^\circ$) expected stimuli should be most active.

Given that the precise nature of mechanisms generating these effects are uncertain, it would be unclear how they would relate to those thought to upweight or downweight expected perceptual representations. Nevertheless, this experiment examined whether perceptual peak shift effects could be observed in action domains, to see whether comparisons can be made to those effects observed in domains outside of action. Experiment 5 was also designed to establish evidence for a perceptual dependent variable that demonstrates learning about action-outcome associations and that would be appropriate for examining neural prediction processes in the fMRI study of Experiment 6. RTs were measured to fulfil this primary aim, but biases in PSEs were also examined to see whether there was any evidence for peak shift effects when action predicts visual outcomes. In contrast to the signal detection paradigm of Cheadle et al. (2015), psychometric functions were modelled to responses from the test session to

determine the extent of perceptual bias in orientation judgements themselves. In the test session, participants responded whether action-outcome orientations appeared closer to a more vertical (closest to 90° and therefore the discrimination boundary) or a more horizontal (45° or 135°) category boundary that represented the extremes of all possible stimulus orientations (see Fig. 4.2.C). These stimulus orientations could be either expected (e.g., clockwise [CW] category) or unexpected (e.g., counter-clockwise [CCW] category) on the basis of action. If expectations bias perception towards exaggerated orientation exemplars of the expected category, akin to typical peak shift effects, then participants will report more often that stimuli are oriented closer to the more horizontal boundary when expected relative to unexpected.

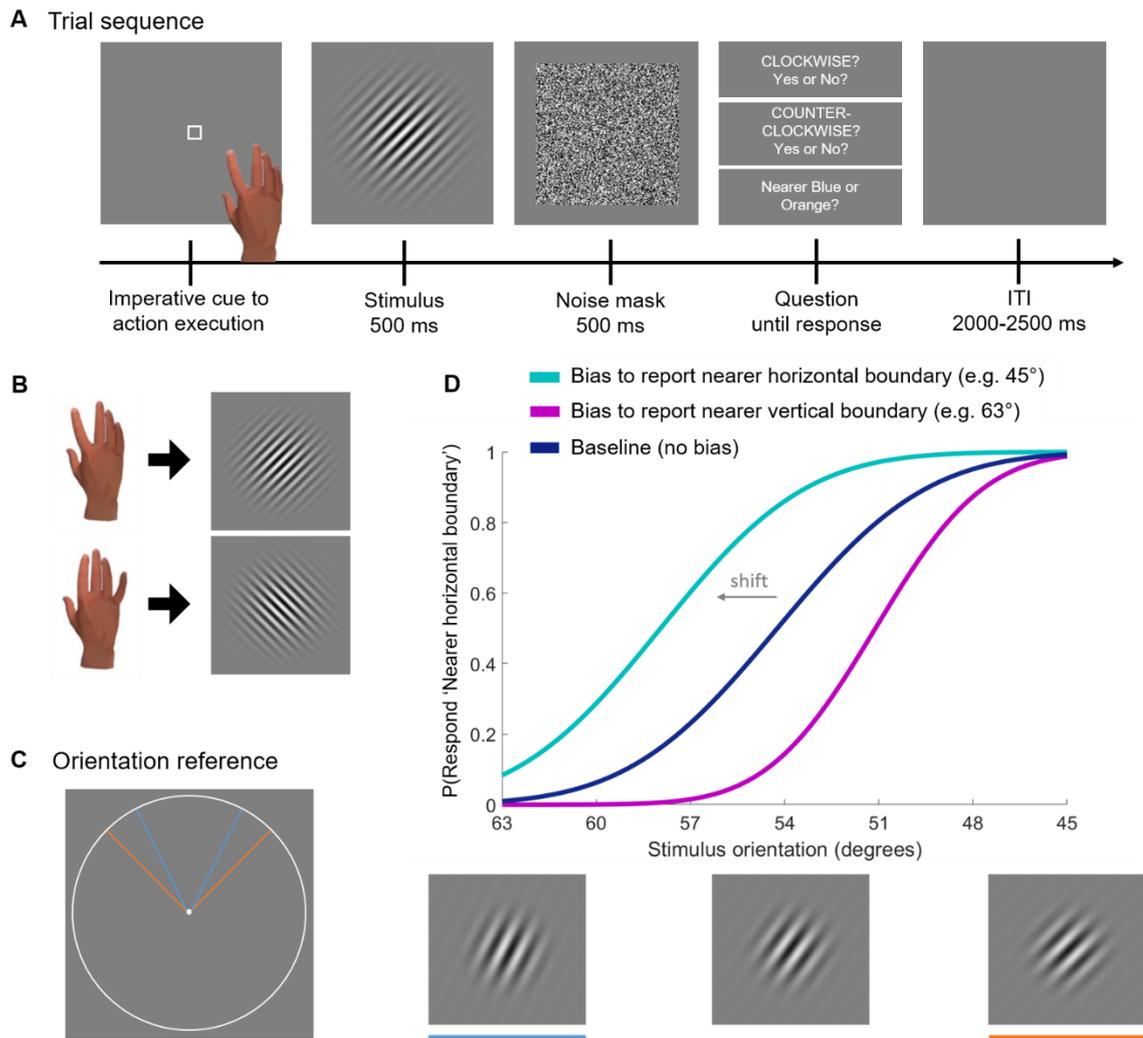


Figure 4.2. (A) A centrally presented visual cue instructed participants to abduct either their index or little finger which produced a CW or CCW grating stimulus. Participants could be asked either whether the stimulus was oriented CW or CCW in part one of the test trials, or whether it was oriented nearer to the blue or orange boundary lines in part two of the test trials. (B) One of two actions was executed on any given trial, whereby an index or little finger abduction perfectly predicted the orientation (CW or CCW) of each stimulus during training. The first part of the test session proceeded with an identical trial sequence except that actions now predicted stimulus orientations with 85% validity. Actions in part 2 of the test session predicted stimulus orientations with 50% validity. (C) In the second part of the test session participants observed the same trial sequence however were instead asked whether the stimulus orientation was nearer to the blue (63° or 117°) or orange (45° or 135°) boundary line, presented every mini-block. (D) A schematic for expected peak shift effects modelled using psychometric functions. A bias towards perceiving extreme exemplars of the expected category (e.g.,

CW stimuli) would lead to a shift of the psychometric curve to the left (turquoise line), reflecting higher probability of responses that stimulus orientations were nearer to the more horizontal boundary (e.g., orange reference line, 45°).

4.1.1. Methods

Participants

Twenty four participants (13 female, mean age = 24.67 years, $SD = 4.34$) were recruited from Birkbeck, University of London and paid a small honorarium for their participation. All participants reported normal or corrected to normal vision. Two participants were excluded from subsequent analyses because psychometric functions could not be modelled effectively in both conditions, resulting in a total sample of 22 participants. The sample size of this experiment was determined to give 80% power of detecting an effect size of 0.55, based upon that observed by Cheadle et al. (2015).

Stimuli

Sinusoidal grating stimuli were created using MATLAB and presented against a grey background using Cogent Graphics. During training, stimuli were presented on a CRT monitor (refresh rate: 80 Hz) at a viewing distance of 45 cm. A Gaussian filter enveloped the grating stimuli to create Gabor patches of 80% Michelson contrast, at 1.5 cycles per degree, and with random spatial phase. Seven Gabor stimuli were generated to appear in both clockwise and counter-clockwise orientations (relative to the hypothetical vertical mid-point i.e. 12 o'clock, 90°), and were each separated by 3° within each category, clockwise (45-63° oriented) and counter-clockwise (117-135° oriented).

Procedure

During training, participants abducted either the index or little finger of the right hand that perfectly predicted the orientation (e.g., CW or CCW) of the stimulus. Participants were instructed to hold down the 'c' and 'm' keys with their respective index and little fingers on a computer keyboard, positioned orthogonally to the screen, until an imperative cue (e.g., square or triangle) indicated which finger to abduct. As soon as a finger lift was executed, the visual stimulus was presented for 500 ms and was followed by a 500 ms noise mask (see Fig. 4.2.A). Stimulus presentation occurred after approximately 12.5 ms (refresh rate at 80 Hz), which resulted in apparent synchrony of stimulus onset with action execution. The stimulus presented was randomly selected from one of seven orientations expected on the basis of the action-outcome mapping. After a variable 300 – 500 ms delay, participants were sometimes asked whether the stimulus was oriented CW and on other trials whether it was oriented CCW. They responded yes or no with their left hand via a handheld keypad. The question (pertaining to CW or CCW) alternated every mini-block of 28 trials. Following response, an ITI of 2000-2500 ms was presented before the start of the next trial.

Twenty-four hours later, participants completed the test trials. The first part of the test session was identical to the training session except that participants' abductions now predicted the stimulus orientation with 85% validity rather than 100% validity. This first part enabled collection of RTs to expected and unexpected stimuli while responses were relevant to the expectation – i.e., relating to orientation. In the second part of the test session, action-outcome contingencies were degraded further such that actions predicted outcomes with 50% validity. Participants were now shown the orientation category boundaries via a reference image (see Figure 4.2.C) which presented boundaries as

either a blue or orange line. In these trials, participants were required to report whether the stimulus was oriented closer to the orange or blue boundary, regardless of whether it appeared clockwise or counter-clockwise (NB: boundary references were only ever referred to by their colours to participants, the terms ‘more vertical’ and ‘more horizontal’ are only used in this thesis for exposition). This reference image was displayed at the start of every mini-block, every 28 trials. All other aspects of this trial structure were identical to the part 1 test and training session.

Participants completed 420 trials in the training session and 476 trials in the test session, which was made up of 196 in part 1 and 280 in part 2. Trial order was randomised and both the specific action-stimulus relationship and the colour order of the orientation reference was counterbalanced across participants. The imperative cue indicating the action to execute on each trial was randomised and this cue-action mapping was reversed half way through each session. Both training and test sessions occurred at the same time on consecutive days. Participants completed 36 practice trials before proceeding to the main trials in session one, and breaks were taken every 28 trials.

Analyses & modelling functions

RT data was collected in the first part of the test session where predictions remained 85% predictive, and median RTs for each participant were calculated for expected and unexpected events. Participant responses were additionally modelled by cumulative Gaussians to estimate psychometric functions (Prins & Kingdom, 2018), using the Palamedes Toolbox in MATLAB. This procedure was performed separately for expected and unexpected trials during the second part of the test session. PSEs were calculated by taking the probability of responses that stimuli were oriented closer to the more horizontal boundary, to describe the point at which participants judged the

stimulus orientation to fall between the two boundaries. Lower values indicate a bias towards perceiving stimuli to be oriented closer to the more horizontal boundary.

4.1.2. Results and Discussion

A paired samples t-test of RTs from the first part of the test session revealed that participants were significantly faster to respond on expected trials ($M = 457.16$ ms, $SD = 249.64$) than unexpected trials ($M = 512.86$ ms, $SD = 276.53$, $t(21) = -2.57$, $p = .018$, $d = -.55$). PSE values of expected and unexpected trials in the second test session were analysed using a paired samples t-test, but revealed no significant differences between expected ($M = 4.51$, $SD = .92$) and unexpected trials ($M = 4.46$, $SD = .94$, $t(21) = 1.31$, $p = .205$, $d = .28$). A Bayesian paired samples t-test (unidirectional, Expected < Unexpected; Cauchy prior width of 0.707) was conducted in JASP (<https://jasp-stats.org/>) to establish the extent of evidence for the null hypothesis, demonstrating moderate evidence, $BF_{0/1} = 9.34$.

The RT effect observed here indicates that participants were able to learn action-outcome associations and therefore learned to anticipate the stimulus orientation on the basis of the action executed. However, actions did not appear to bias PSEs in a manner that would reflect a peak shift effect. Even numerically, perception was not biased towards perceiving the exaggerated exemplar of the learned category (e.g., towards the more horizontal boundary) when it was expected relative to unexpected.

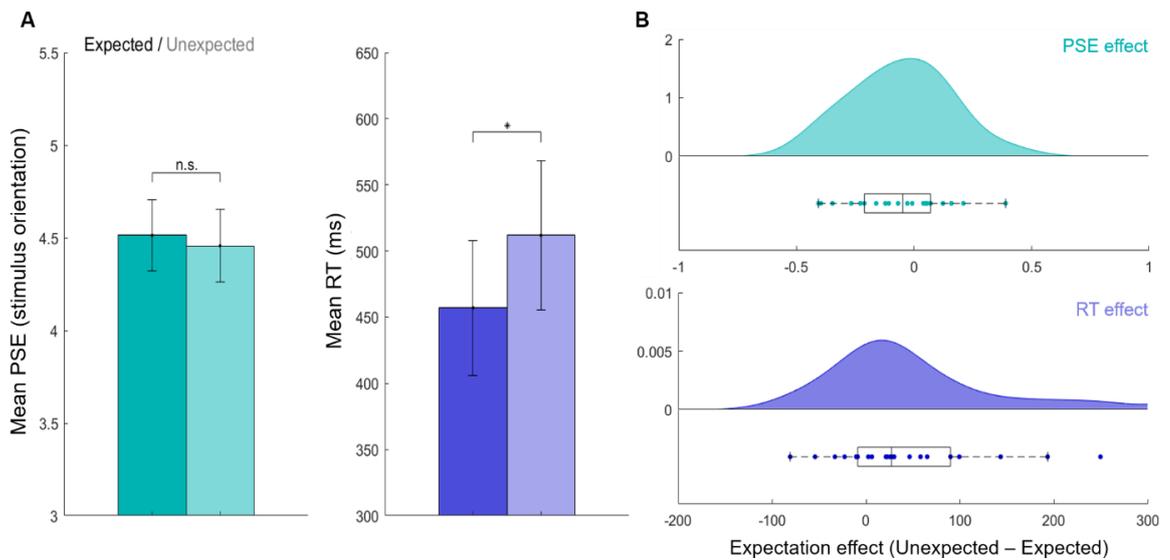


Figure 4.3. (A) Left panel: Mean PSEs (\pm SEM) for expected and unexpected trials during the second part of the test session, which were not significantly different. PSEs were calculated on the probability of responses that stimulus orientations were ‘more horizontal’. PSE y-axis units index stimulus orientation, ranging from the most vertical (stimulus level 1) to the most horizontal (stimulus level 7) boundaries, averaged across both CW and CCW orientations. Lower values indicate a bias to perceive stimulus orientations closer to the more horizontal boundary. Right panel: Mean RT (\pm SEM) expectation effect during the first part of the test session. $*p < 0.05$. (B) Expectation effects (Unexpected – Expected) plotted with raincloud plots (Allen et al., 2019) displaying probability density estimates (upper) and box and scatter plots (lower), for PSEs (top) and RTs (bottom). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=30). Positive PSE expectation effect values indicate perceived orientations closer to the more horizontal boundary for expected events relative to unexpected events, while positive RT effects reflect faster responses on expected trials.

The results are therefore inconsistent with previous peak shift findings and the visual prediction-related effects observed in Cheadle et al. (2015), that expectations bias perception towards exaggerated exemplars of that expected category. However, there are a number of noticeable differences between the typical studies demonstrating these effects and the present experiment. Specifically, much of this literature does not focus

on demonstrating predictive influences, but rather trains participants to discriminate between two categories of stimuli in comparison, while in the present experiment categories of stimuli were associated with action. Furthermore, it is possible that some peak shift effects observed in the literature are not perceptual in nature and could instead be explained by post-perceptual decisional biases. For instance, studies have found that when discrimination boundary lines are not presented during response, or reproduction, such biases are eliminated (Zamboni et al., 2016). Additionally, studies that measure participants' reproduction of oriented lines or dot-motion (e.g., Jazayeri & Movshon, 2007) without presenting the boundary reference simultaneously may render responses that rely too much upon working memory components, reflecting post-perceptual decisional biases rather than perceptual biases (Fritsche & de Lange, 2019). Fritsche and de Lange (2019) further tested whether peak shift, or 'reference repulsion' effects, could be explained purely by perceptual biasing when reducing the working memory load and requiring judgements about stimulus orientation in direct comparison to a reference on screen. This experiment revealed a repulsion bias for stimuli oriented further away from the discrimination boundary but demonstrated a pattern that was more akin to classical adaptation effects – concluding that typical reference repulsion effects may indeed be better explained by post-perceptual decision biases or influences of working memory (Fritsche & de Lange, 2019).

Such findings could explain why this experiment failed to replicate previous peak shift effects, especially when the present experiment adopted different measurements of perceptual bias (e.g., PSEs) than have been previously examined in the peak shift literature. However, this experiment did observe RT benefits for expected outcomes, which most importantly suggests that participants had learned to expect orientations associated with action. We cannot know from RT measurements that these benefits are

produced by perceptual biasing given that the response question was also predictable in this experiment, and therefore anticipation of responses on expected trials could have occurred regardless of actual sensory evidence. Nevertheless, such RT biases are consistent with a variety of other experiments adopting manipulations of prediction according to statistical regularities outside of action and demonstrating faster RTs for expected than unexpected events (Kok et al., 2014; Richter et al., 2018; Richter & de Lange, 2019). Such RT benefits have even been shown to correlate with neural expectation effects in V1 (Richter & de Lange, 2019) – suggesting that these findings could indeed plausibly be produced by perceptual biases. It should also be pointed out that PSE and RT effects were measured in different parts of the test session, and actions therefore had stronger predictive validity of the stimulus when RTs (85%) were measured compared to when PSEs (50%) were measured and therefore comparability is limited.

A perhaps more plausible explanation for the lack of perceptual biasing in the PSE measure relates to literature suggesting that perception might instead be biased towards the middle or average representation of a learned category. A study by Ester et al. (2020) trained participants to discriminate between two categories of oriented stimuli and recorded brain activity in visual areas during categorisation – finding that reconstructed representations of stimuli yet to be categorised were biased towards the centre exemplar of the appropriate category. This effect might be akin to a well-known phenomenon in perceptual and memory research, in which it is common for observers to bias estimates of categorised stimuli towards the most typical or centre of a stimulus set (Crawford et al., 2000). These biases are thought to be adaptive under a Bayesian process that improves estimate accuracy and helps to resolve representation uncertainty, especially when a category contains many stimulus exemplars (Duffy et al., 2010;

Huttenlocher et al., 2000). Such a bias may be present in the current data, given that there were seven possible stimulus orientations for each category, thereby resulting in a biasing of perception towards the most typical of expected stimuli to form accurate representations of the category under uncertainty. The present set-up, however, was not optimal for detecting such a bias considering that participants were required to respond whether the stimulus orientation was closer to the more vertical or horizontal boundaries, and therefore cannot distinguish a bias towards the middle of the category from baseline performance when measured in this way.

It is unclear whether such an interpretation would be consistent with predictive upweighting or downweighting mechanisms since this phenomenon is not often considered in terms of predictive processing. Furthermore, not enough is known about whether peak shift effects reflect perceptual biasing and therefore whether they could operate according to either of the proposed predictive mechanisms. Nevertheless, the primary question of importance from this experiment was whether a measure could be obtained that demonstrated learning about action-outcome pairings in this paradigm. Indeed, RT expectation effects were observed when actions predicted oriented gratings – akin to in other neural paradigms outside of action measuring similar processes – therefore demonstrating that this provides an appropriate paradigm for investigating the neural influences of prediction during action.

4.2. Experiment 6

Neural support for the cancellation account demonstrates reduced activity in sensory brain regions for self-produced action-outcomes relative to externally generated outcomes. For example, self-generated tactile and visual events elicit a reduced neural response in somatosensory and visual cortices (respectively) relative to externally generated events (Blakemore et al., 2001; Kilteni & Ehrsson, 2020a; Shergill et al., 2013; Stanley & Miall, 2007; Wolpe et al., 2018). Additionally, studies outside of action domains have also demonstrated suppression of activity associated with expected relative to unexpected events (Alink et al., 2010; Han et al., 2019; Kok et al., 2012; Richter et al., 2018; Todorovic & De Lange, 2012). Interestingly, the study by Kok and colleagues (2012) demonstrated that expectation suppression may not necessarily reflect downweighting mechanisms but instead upweighting ‘sharpening’ mechanisms – whereby sensory populations tuned to predicted events are activated, yet populations tuned to unpredicted or irrelevant events are suppressed (de Lange et al., 2018; Kok, Jehee, et al., 2012; Press et al., 2020b; Summerfield & de Lange, 2014).

Recent evidence has additionally suggested that sharpening mechanisms may also influence sensory processing during action. Yon and colleagues (2018) used MVPA to show that visual hand movements that were congruent with participants’ own actions were more readily decoded from patterns of visual activity than unexpected events. Additionally, univariate stimulus-specific analyses in these same areas demonstrated that expectation suppression (for congruent relative to incongruent outcomes) was only present in voxels that were tuned *away* from the presented stimulus. For example, when participants observed an index finger movement, visual activity was lower when their own executed action was congruent with this visual input, relative to when their action was incongruent – yet this pattern was only observed in voxels that were not tuned to

index fingers but to little finger movements. These findings mirror those of Kok et al.'s (2012) study that showed superior pattern classification of expected events predicted by an auditory cue relative to unexpected events, and are consistent with theories suggesting that expectations sharpen sensory representations – which are more consistent with upweighting perceptual accounts than downweighting accounts. Therefore, these findings might suggest that prediction sharpens sensory processing in early visual areas in a similar fashion regardless of whether the predictive cue is action or another sensory event.

However, although the study by Yon and colleagues (2018) examined the relationship between action and congruent vs. incongruent outcomes – which likely reflects natural predictions learned over the life span e.g., lifting my index finger results in visual input of index finger lifting up – the predictive nature of these effects is uncertain considering that statistical probabilities were not directly manipulated. Experiment 6 therefore provides a conceptual replication of the studies by Yon et al. (2018) and Kok et al. (2012) to examine the precise nature of predictive mechanisms when visual events are trained via statistical learning to be expected or unexpected action-outcomes.

Participants learned during training that one of two actions (finger abductions) could predict one of two stimulus orientations (CW or CCW), similarly to in Experiment 5. During scanning, the relationship between action and stimulus orientation was degraded such that expected and unexpected events were equally likely to occur.

If sensorimotor prediction influences sensory processing in the same way as has been observed in previous literature (Kok, Jehee, et al., 2012; Yon et al., 2018), it would be hypothesised that expected events will be decoded using MVPA with superior accuracy than unexpected events and this pattern will be accompanied by univariate expectation

suppression only in voxels tuned away from the presented stimulus. If however, action predictions influence sensory processing in a distinct manner, then an alternative pattern is likely to be observed, and one possibility for such an alternative might be neural downweighting patterns.

4.2.1. Methods and Analyses

Participants

Twenty-one new participants were recruited (13 female, mean age = 27.91 years, $SD = 5.50$), from Birkbeck College, University of London and University College London (UCL) and paid a small honorarium for participation. All participants reported normal or corrected to normal vision and had no history of psychiatric or neurological illness. One person was excluded for obtaining movement parameters which exceeded set thresholds (sudden movements larger than 3 mm throughout runs), resulting in a final sample of 20 participants. The experiment was approved by both Birkbeck and UCL ethics committees. The sample size for this experiment was selected on the basis of previous studies (Kok, Jehee, et al., 2012; Yon et al., 2018) that performed similar analyses and observed reliable univariate and multivariate effects.

Stimuli

Differences from the sinusoidal grating stimuli used in Experiment 5 are described in the following section. During training, stimuli were presented on a Dell Laptop via an LCD screen (resolution: 1280x1024; refresh rate: 60 Hz) at a viewing distance of 45 cm, and during scanning on a LCD monitor (resolution: 1280x1024; refresh rate: 60 Hz) at a viewing distance of 77 cm. In both cases, stimuli were viewed at 15 degrees of visual angle. Stimuli were now presented in an annulus around a fixation cross in the middle of the screen (see Fig. 4.4.A). Two stimulus orientations were generated to

appear in both clockwise and counter-clockwise orientations (relative to the hypothetical vertical mid-point e.g., 12 o'clock, 90°), at 45° (clockwise oriented) and 135° (counter-clockwise oriented) degrees.

Procedure

Differences in procedure from Experiment 5 are hereby described. Each trial started with a white fixation cross, which remained on screen until participants' response. On stimulus trials, once the appropriate action was executed the imperative cue was replaced with an oriented stimulus for 500 ms. Participants were required to respond to the question screen within 1500 ms and the next trial started after a variable ITI of 2-3 s (in training) or 2-6 s (in scanning). Omission trials were also included, where the imperative disappeared leaving just the fixation cross on the grey background. The data from these trials are not presented in this thesis.

The experiment consisted of training and test sessions split over two consecutive days and occurring at the same time on both days. In the test session, participants lay in the scanner with their left and right hands on separate MR-compatible button boxes. The right hand button box was positioned parallel to body midline, such that the index finger was positioned superior to the little finger. Participants depressed two buttons on the right hand box with their index and little finger for the duration of runs, except during action execution or during breaks. The left hand button box was positioned inferior to the right hand box and placed by the proximal part of their left leg, and participants placed their left thumb between two keys for response. Actions produced expected or unexpected stimulus orientations (33.3% validity respectively) as well as omission trials (final 33.3%) with equal likelihood. The experiment was carried out in ten scanning blocks of 36 trials each.

In the training session outside of the scanner, participants completed a 90-minute task in which actions were only ever superseded by stimuli (and never omissions) and were perfectly predictive of the stimulus orientation (100% contingency – e.g., index finger abduction always resulted in presentation of a CCW oriented stimulus). The set-up was otherwise identical in the scanning session except that actions were executed via MR compatible button boxes positioned orthogonal to the laptop screen. A short but identical training session (~15 mins) was presented to participants immediately before they went into the scanner in the test session similarly using the MR button boxes.

Participants completed 360 main experimental trials in the scanning session and 600 training trials in the first session. The short training session before entering the scanner consisted of 120 trials. During training sessions, the response question alternated every mini-block, but was randomised on each trial during scanning to unconfound the response with stimulus orientation in later analyses. Imperative cue order and trial order were randomised within blocks and the specific action-stimulus relationship was counterbalanced across participants. The imperative cue-action mapping was also counterbalanced and reversed halfway through each session (e.g., at the beginning of the sixth block) to unconfound potential influences of cue-outcome learning.

Participants completed 36 practice trials before proceeding to the main training trials in session one, and breaks were taken every 36 trials in all sessions.

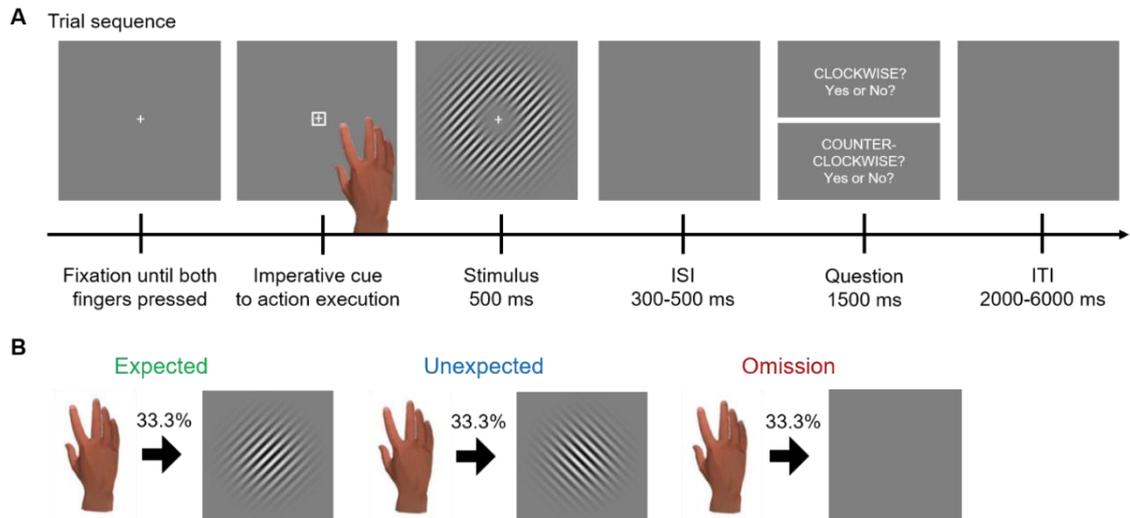


Figure 4.4. Experimental paradigm. (A) A centrally presented visual cue instructed participants to abduct either their index or little finger. Each finger abduction predicted an oriented stimulus and participants were required to respond (yes/no) to whether the stimulus was clockwise or counter-clockwise. In training, actions perfectly predicted the stimulus orientation that was presented. (B) In the scanning session 24hrs later, participants completed the same task but actions now predicted the oriented stimulus with 33.3% validity, presenting expected and unexpected stimuli 33.3% of the time each and on the remaining trials the stimulus was subsequently omitted.

Behavioural analyses

RT data was collected for responses to expected and unexpected stimuli in the test session and median RTs were calculated across all trials separately for each condition, for each participant. Similarly, the proportion of correct responses was analysed for expected and unexpected conditions for each participant.

Image Acquisition

Images were acquired using a 3T Prisma MRI scanner (Siemens Healthcare GmbH, Erlangen, Germany) using a 32-channel head coil at the Wellcome Centre for Human Neuroimaging (WCHN), UCL. Functional images were acquired using an echo planar imaging (EPI) sequence (ascending slice acquisition, TR = 3.36 s, TE = 3.29 s, 48

slices, voxel size: 3 x 3 x 3 mm), and anatomical images were acquired using a magnetisation-prepared rapid gradient-echo (MPRAGE) sequence (voxel size: 1 x 1 x 1 mm).

fMRI data preprocessing and ROI definition

The analysis techniques performed here are almost identical to those presented in Yon et al., (2018), except for the region of interest (ROI) definition. Preprocessing of the images was conducted in SPM12. The first six volumes of data were discarded from each scanning run for each participant to allow for T1 equilibration. Functional images were spatially realigned to the mean image to yield head movement parameters, and were temporally-realigned to the 24th (middle) slice. The realigned functional images were then coregistered to participants' structural scan and no normalisation or smoothing procedures were performed due to the nature of these analyses.

Freesurfer (<http://surfer.nmr.mgh.harvard.edu/>) was used to define the ROI, V1, based on anatomical landmarks and by computing the probability of inclusion for each voxel using pre-established methods (Hinds et al., 2008) for each participant (for example see Fig. 4.5.A). All voxels that were estimated to have at least 0.5 probability of being present in V1 were selected in ROIs for further analysis.

GLM specification

A general linear model (GLM) was first specified in SPM12 for each of the 4 regressors of interest and included movement parameters as 6 nuisance regressors. These regressors of interest modelled each stimulus orientation that was presented (CW, CCW) in each experimental condition (expected and unexpected) for each scanning run. Regressors were modelled to the onset of the presented stimulus and convolved with the

canonical haemodynamic response function. This GLM generated 4 beta images (one for each regressor of interest) in each of the 10 scanning runs that were used for subsequent analyses.

Multivariate Pattern Analysis

MVPA analyses were conducted using the TDT toolbox (Hebart et al., 2015) and involved training linear Support Vector Machines (SVMs) to discriminate between clockwise and counter-clockwise orientations in the BOLD activation patterns from trials where stimuli were presented (not omissions) across voxels in V1. A leave-one-out cross-validation procedure was adopted which included 10 decoding steps, where different SVM classifiers were trained and tested on the beta images separately for each experimental condition (expected and unexpected). The GLM produced 10 beta images (one for each scanning run) for each stimulus orientation (CW, CCW), resulting in 20 beta images for each experimental condition. Each decoding step used 18 beta images (CW and CCW) from nine scanning runs for classifier training to estimate a linear discriminant function separating CW and CCW stimulus orientations. This was then applied to the two beta images from the remaining tenth run to classify them as CW or CCW orientations. Accuracy of the SVM's was then calculated as an average of the proportion of correctly classified images across all decoding steps.

Univariate analysis

This analysis investigated how patterns of stimulus-specific univariate BOLD activity differed according to expectation in V1. Taking the unnormalised and unsmoothed beta images of presented stimuli as used in the decoding analysis, activity for CW and CCW orientations was contrasted in SPM12 for each participant. This yielded a t-contrast map containing orientation preference information for each voxel for each participant –

where positive t-statistics ($t > 0$) indicated a preference for CW oriented stimuli and negative values ($t < 0$) indicated preference for CCW oriented stimuli.

Univariate BOLD signals (beta values) were extracted from voxels according to their assigned orientation preference, and in each of the experimental conditions (expected and unexpected) to calculate signals that correspond to factors of interest, Expectation and Preference. For example, if a voxel was classified as ‘CW preferring’, and a CW grating was presented on an expected trial, the univariate signal for this voxel would be categorised under ‘Expected Preferred’, whereas signals in voxels classified as ‘CCW preferring’ on the same trials would be categorised as ‘Expected Non-Preferred’. These signals were extracted from each voxel in V1 and analysed with a repeated measures ANOVA with the factors Expectation (expected, unexpected) and Preference (preferred orientation, non-preferred orientation).

4.2.2. Results and Discussion

Behavioural analysis

RT analyses revealed no significant difference between response times for expected ($M = 668.98$ ms, $SD = 84.74$) and unexpected ($M = 695.83$ ms, $SD = 120.54$) trials ($t(19) = -1.25$, $p = .228$, $d = -0.28$). Participants were significantly more accurate for expected ($M = .98$, $SD = .03$) trials than unexpected ($M = .96$, $SD = .03$) trials ($t(19) = 5.82$, $p < .001$, $d = 1.30$).

MVPA

The accuracy of SVMs over all of the decoding steps were calculated as a proportion of correctly classified beta images as either CW or CCW oriented, and represent percentage accuracy above chance in each condition. Comparison of these decoding

accuracies revealed a significant effect of Expectation ($t(19) = 3.27, p = .004, d = .73$) such that there was superior decoding of stimulus orientation when the outcome was expected ($M = 9.25, SD = 7.66$) compared to unexpected ($M = 2.75, SD = 9.52$; see Fig. 4.5.C), suggesting that enhanced stimulus orientation information could be read out from the BOLD signal when stimuli were expected.

Univariate analysis

To examine evidence for expectation suppression in voxels that were either tuned towards or away from the presented stimulus, a 2 x 2 repeated measures ANOVA was conducted on stimulus-specific univariate BOLD signals. This analysis revealed a significant main effect of Expectation ($F(1,19) = 17.75, p < .001, \eta_p^2 = .48$), Preference ($F(1,19) = 1094.58, p < .001, \eta_p^2 = .98$) and a significant interaction between Expectation and Preference ($F(1,19) = 41.14, p < .001, \eta_p^2 = .68$). This interaction was driven by lower activity in voxels tuned away from the presented stimulus for expected ($M = 6.62, SD = 1.15$) relative to unexpected ($M = 6.88, SD = 1.11$) events ($t(19) = -8.01, p < .001, d = -1.79$), but no difference in voxels tuned towards the presented stimulus for expected ($M = 8.58, SD = 1.15$) and unexpected ($M = 8.62, SD = 1.08$) events ($t(19) = -.92, p = .371, d = -.21$; see Fig. 4.5.D). This pattern therefore indicates that action predictions evoke expectation suppression only in voxels tuned away from, rather than tuned to, the presented stimulus.

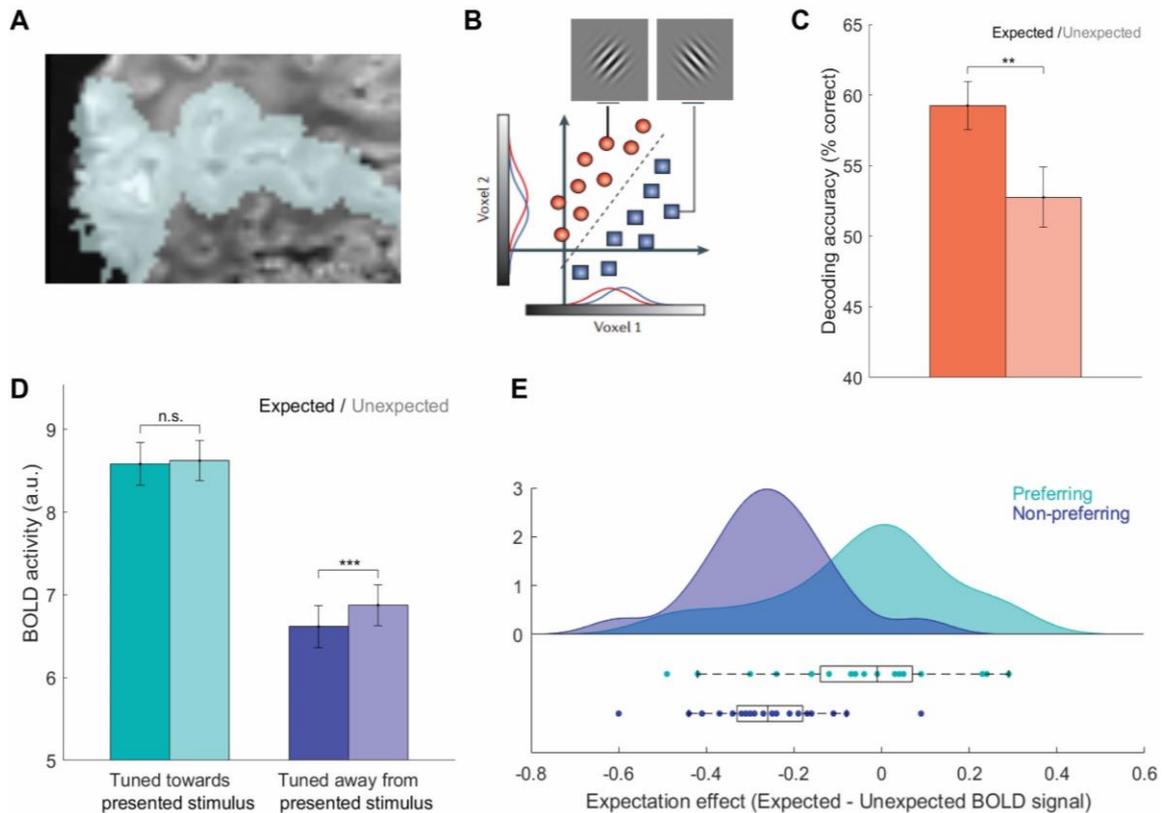


Figure 4.5. (A) Sagittal view of an example mean functional image of the calcarine sulcus and occipital lobe overlaid with the estimated V1 ROI (in white). (B) A schematic representing classification of CW and CCW activity patterns across voxels using MVPA, adapted from Haynes and Rees (2006). (C) Mean decoding accuracy (\pm SEM) for expected and unexpected events, showing superior classification of stimulus orientation when stimuli were expected. (D) Mean BOLD activity (\pm SEM) for voxels tuned towards and away from the presented stimulus separately for expected and unexpected events. Expectation suppression is only observed in voxels tuned away from the presented stimulus (** $p < .01$, *** $p < .001$). (E) Expectation effects (Expected BOLD – Unexpected BOLD signals) were calculated for voxels tuned towards [teal], and away from [blue] the presented stimulus and are plotted with raincloud plots (Allen et al., 2019) displaying probability density estimates (upper) and box and scatter plots (lower). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant (N=20). Negative effects indicate lower BOLD signal for expected events relative to unexpected events.

These results demonstrate that stimulus orientations could be decoded from activation patterns in V1 with superior accuracy on expected relative to unexpected trials – signalling enhanced information about the stimulus when it was expected. Univariate analyses demonstrated that expectation suppression was found only in voxels that were tuned away from the presented stimulus – consistent with hypotheses from sharpening theories that expected activity is only suppressed in populations not tuned to the expected stimulus. These results are in line with previously discussed findings demonstrating that expected representations are activated, while activity in populations tuned away from expected stimulus features is suppressed (Kok, Jehee, et al., 2012; Yon et al., 2018).

Experiment 6 explicitly manipulated the statistical probability of expected and unexpected events and demonstrates comparable effects to Yon et al.'s (2018) study that examined congruency between action-outcomes. These similarities suggest that both manipulations likely reflect underlying sharpening processes operating in similar fashions, given that congruent associations are likely learned throughout the life span. Taken together, these findings suggest that predictions formed on the basis of statistical associations with action also operate according to sharpening mechanisms, similarly to prediction mechanisms outside of action adopting comparable manipulations of statistical learning.

The results are however inconsistent with the conclusion from an array of studies in action demonstrating suppressed activity evoked by predictable consequences of action, that are frequently taken to support downweighting cancellation theories (Blakemore et al., 1998; Kilteni & Ehrsson, 2020a; Kontaris et al., 2009; Shergill et al., 2013; Stanley & Miall, 2007). One explanation that might account for differences across domains

relates to assumptions that tactile attenuation during action is dependent upon somatosensory-cerebellar connectivity (Blakemore et al., 1998; Kilteni & Ehrsson, 2020a), in contrast with hippocampal mediation of prediction in visual processing (Kok et al., 2020; Kok & Turk-Browne, 2018). However, the present findings indicate the possibility that such conclusions are unlikely because neural suppression findings in the literature may instead be underpinned by sharpening mechanisms. Furthermore, Experiment 6 demonstrates that prediction mechanisms in visual regions likely operate in the same fashion regardless of whether sensory outcomes can be predicted by action or other sensory events, and similar conclusions regarding influences on tactile perception presented earlier in this thesis perhaps signal that somatosensory neural mechanisms may also operate similarly. However, further neural work presenting tactile outcomes would be necessary to examine this claim.

4.3. General Discussion

This chapter examined how action predictions influence visual perception and sensory processing in V1 to determine whether expectations sharpen stimulus-specific representations in line with upweighting theories, or whether representations of expected action outcomes are suppressed as proposed by downweighting theories. Experiment 5 first provided a pilot experiment for the fMRI paradigm of Experiment 6, to examine evidence for any behavioural effects indicating that participants could learn to expect outcomes on the basis of action. It indeed found RT benefits for expected events, demonstrating that participants could learn action-outcome mappings in this paradigm. It also examined a hypothesis that expectations bias perception towards exaggerated exemplars of expected stimuli but found no evidence to support this claim. Experiment 6 therefore used the same paradigm to investigate neural prediction mechanisms during action and found evidence in favour of the operation of sharpening

mechanisms. Specifically, enhanced information about stimulus orientation could be decoded from expected signals, while expectation suppression was observed only in voxels that were tuned away from the presented stimulus. These results are in line with predictive upweighting theories of perceptual processing, proposing that we use prior knowledge or expectations about the statistics of our environment to enhance perception – thereby aiding veridicality of our experiences (Press et al., 2020b). Importantly, these findings demonstrate that action prediction mechanisms appear to operate in qualitatively similar ways to neural mechanisms outside of action.

One feature of this design involved degrading the contingency so dramatically that ‘expected’ events were in fact only presented a third of the time during the scanning session. In theory, expected orientations were relatively unlikely to occur (33% of the time) in combination with unexpected stimuli and omissions (remaining 67%).

Furthermore, on stimulus trials, expected and unexpected events were equally likely to occur, but regardless the evidence demonstrates highlighted processing of expected orientations on the basis of previously learned associations. This is interesting because there have been some recent suggestions that effects of expectation need to be compared with ‘neutral’ conditions – in which no expectations are formed about the resulting stimuli – to accurately assess whether effects reflect an increase in unexpected activity or a decrease in expected activity (Feuerriegel et al., 2020). Such propositions have suggested that contingencies of 50% (equal probability) would be appropriate to achieve these aims, such that presenting two possible outcomes with equal probability will likely render no expectations about upcoming input. The findings of Experiment 6 highlight how in fact pre-learned associations still greatly influenced perceptual processing even when they become less statistically probable at test – perhaps suggesting that these pre-learned associations are fairly robust. Nevertheless, it could be

plausible to hypothesise that the strength of predictive effects should be modulated by the extent of contingency that remains between associated events, and therefore larger RT effects observed in Experiment 5 than 6 could be explained by the remaining predictive strength of associations, given that actions in the former still predicted events with 85% probability.

Notably, these findings are inconsistent with downweighting theories that propose that perceptual processing of expected outcomes is suppressed in order to highlight more informative unexpected events. Studies providing support for such accounts in action rarely demonstrate that self-generated outcomes are predictable according to statistical probabilities within the environment, and instead compare perceptual processing of such events (assumed to be predicted by action) with externally generated events (Blakemore et al., 1998; Kilteni & Ehrsson, 2020a; Shergill et al., 2013; Stanley & Miall, 2007). It is therefore unclear whether the underlying mechanisms generating these suppression effects in action are predictive when explicit predictive manipulations have not been performed.

Intriguingly, some recent evidence outside of action domains that adopts statistical learning paradigms to investigate underlying prediction mechanisms in sensory regions, in fact demonstrates support for stimulus-specific downweighting, or otherwise termed ‘dampening’ accounts (Han et al., 2019; Richter et al., 2018; Richter & de Lange, 2019). For example, in two studies by Richter and colleagues (2018; 2019), participants were trained to learn the relationship between several pairs of objects, where one object could predict the identity of a second object. The authors revealed that expectation suppression was present across V1 and Lateral Occipital Cortex (LOC) for predictable objects, however this scaled positively with image preference such that the largest

effects of expectation suppression were found in voxels that were tuned towards the presented image – consistent with downweighting rather than upweighting, sharpening, accounts. These findings suggest that the neural picture may be more complex than assumed until recently, even outside of action domains. One difference between Richter and colleagues’ paradigms and those demonstrating support for sharpening accounts is that there are a large number of associations between object pairs being learned and updated across these experiments, in comparison to just two in Experiment 6, Yon et al. (2018) and Kok et al. (2012). This set up will generate a scenario whereby each unexpected event is relatively more unexpected than those in the present study, potentially requiring more deviance detection and model updating. Under the opposing process theory this difference could be crucial because it will increase the likelihood of operation of the later process which perceptually highlights unexpected events (Press et al., 2020b). However, given there are a range of differences between studies, any conclusions concerning these specific differences are tentative – though these findings highlight the possibility that both predictive upweighting and downweighting mechanisms might operate under different contexts. Nevertheless, the important conclusion from Experiment 6 is that when you employ similar manipulations in action that have demonstrated predictive sharpening mechanisms outside of action, evidence for sharpening is also found within action, suggesting that similar mechanisms likely underlie prediction in both domains.

Chapter 5: The laminar profile of surprise

Chapter 4 demonstrated that sensorimotor prediction likely generates sharpening of perceptual representations in early visual regions. These findings suggest that action predictions influence perceptual processing in a similar fashion to predictions outside of action, that primarily operate according to upweighting principles (de Lange et al., 2018; Kok, Jehee, et al., 2012). However, a conclusion that all predictive influences on prediction are generated by upweighting mechanisms would be premature before further investigations that attempt to examine how stimulus-specific downweighting findings – e.g., recently demonstrated on the basis of statistically probable expected outcomes (Han et al., 2019; Richter et al., 2018; Richter & de Lange, 2019) – could be reconciled with upweighting findings.

A possible resolution to the inconsistencies across and within domains could come from findings reporting opposite influences of prediction in different cortical regions, and could support a hypothesis that both up- and downweighting processes operate but at different points of the cortical hierarchy, and perhaps with different functions. For example, while studies examining predictive mechanisms in early sensory regions have observed expectation suppression (Han et al., 2019; Richter et al., 2018; Richter & de Lange, 2019), studies examining processing in frontal regions e.g., prefrontal cortex (PFC) have demonstrated enhanced signals associated with prediction (González-García & He, 2021; Rahnev et al., 2011). Conversely in action, activity in frontal regions such as the superior and middle frontal gyri (SFG and MFG) is reported to be suppressed for predictable self-generated sounds, while in primary auditory cortex, superior temporal gyrus (STG), self-generated sounds elicit enhanced activity relative to passive listening of the same speech (Reznik et al., 2014). Other similar action studies have shown

differences between neighbouring regions, with active speech eliciting attenuated BOLD signal in STG but enhanced signal in superior temporal sulcus (STS) relative to passive listening (Christoffels et al., 2007). However, as previously noted in this thesis, conclusions about such inconsistencies would be difficult when comparisons are made between active and passive perception in these latter examples, because pure predictive influences cannot be determined.

Moreover, a study outside of action has suggested that both expectation suppression and enhancement effects found in early sensory and frontal regions, respectively, may in fact both reflect representation-specific sharpening mechanisms. Specifically, González-García and He (2021) found that the default mode network (DMN) exhibits enhanced activity for ambiguous patterns that can be predicted on the basis of prior knowledge, relative to the same patterns that cannot be predicted, yet activity was suppressed for the same predictable patterns in visual areas (V1-4). Interestingly, both of these effects were revealed to reflect sharpened stimulus-specific representations when effects were examined across voxel-wise image preference. Specifically, expectation suppression in visual areas was the largest in voxels least selective to presented images, while expectation enhancement in frontal DMN areas was largest in the most selective voxels.

Although providing some preliminary evidence that predictive processes perhaps operate in different manners across different regions in the cortical hierarchy, these studies are unable provide a unified account of how opposing up- and downweighting mechanisms might be implemented in different regions to explain the discrepancies between the reported sharpening and cancellation findings. However, considering that there are also discrepant findings observed within regions (e.g., in V1; Kok et al., 2017; Kok, Jehee, et al., 2012; Richter et al., 2018; Richter & de Lange, 2019), differences

may instead be better reconciled by a hypothesis that both processes occur within the same sensory regions but at different cortical depths. The resulting effects observed in particular regions across studies may therefore depend on the particular design and process that is being relied on, determining which particular processes ‘wins’ in determining the observed neural patterns.

Such a conclusion would be supported by a key tenet of predictive coding theories, that prediction (P) and prediction errors (PE) are represented in distinct units within a cortical layer and with different functional roles (Bastos et al., 2012; Friston, 2005). For example in V1, predictions conveyed through feedback connections are thought to reside in the deep layers of cortical regions while prediction error signals are passed up via feedforward connections and thought to reside primarily in superficial layers (Bastos et al., 2012). Such a hypothesis could suggest that upweighting processes that more highly weight expected outcomes are represented in the activity of prediction neuronal populations in deep layers, while downweighting processes that highlight unexpected events are represented in the activity of prediction error populations in superficial layers. More specifically, expectation suppression may therefore be a product of greater activity in superficial layers from PE units carrying feedforward information about unexpected events, while P units carrying feedback information about expectations in deep layers produce effects akin to sharpening.

There is some evidence from monkey studies to suggest that prediction processes may operate differently across layers within cortical regions. This work often invasively records neural firing rates in sensory areas when expected and unexpected stimuli are presented, and has demonstrated increased neural firing rates in superficial and deep layers of V1 during top-down figure filling in the absence of bottom-up input (e.g.,

Lee & Nguyen, 2001). In humans, the relatively recent development of non-invasive ultra-high resolution 7T MRI methods has allowed examination of human laminar activation patterns because voxel sizes smaller than the cortical thickness can be acquired to measure differences across cortical layers. Similarly, studies using these methods in humans have provided evidence for prediction related activity in deep cortical layers of V1 in figure ground segregation tasks (Kok et al., 2014) as well as when events are expected but subsequently omitted (Aitken et al., 2020). However, there is a surprising absence of studies designed to investigate the claim that prediction error related activity is represented in superficial layers.

This chapter therefore aims to examine the claim that prediction and prediction error encoding predominates in different cortical layers, and therefore whether such a distinction could help to reconcile the range of paradoxical findings reporting evidence for both upweighting and downweighting predictive processes. Experiment 7 first presents a psychophysical study that examines the extent to which different contingencies between action and outcomes influence the magnitude of expectation effects on perception. This paradigm is then adapted in Experiment 8 for a 7T fMRI study to examine the laminar profile of expected and unexpected events in visual regions.

5.1. Experiment 7

Influences of prediction on perceptual processing have been measured in a variety of ways, from top-down figure filling-in (Kok, Bains, Van Mourik, et al., 2016; Lee & Nguyen, 2001; Self et al., 2013; Smith & Muckli, 2010) and disambiguation of noisy image features (Flounders et al., 2019; González-García & He, 2021) to training of associations between sensory events through statistical probabilities (Aitken et al., 2020; Kok et al., 2014, 2017; Kok, Jehee, et al., 2012; Kok & Turk-Browne, 2018; Richter & de Lange, 2019). One of the key principles of this thesis is to determine how perceptual processing is influenced by sensorimotor predictions manipulated according to the latter – relative to statistical probabilities of sensory events. In studies using such statistical manipulations, sensory cues often predict visual outcomes with 75% validity and demonstrate behaviourally faster responses to expected than unexpected trials (Kok et al., 2014; Kok, Rahnev, et al., 2012). Other studies have demonstrated that such RT benefits are correlated with neural expectation effects (Richter & de Lange, 2019). Comparable RT benefits were found in Experiment 5 when the contingency between action and outcome remained high (85%), but this effect was abolished in Experiment 6 when the contingency was low (33%). Regardless of the lack of behavioural effect in the latter, evidence for the operation of neural prediction mechanisms was still observed. However, given that previous findings indicate RT benefits are correlated with neural expectation effects, Experiment 7 was designed to maximise a behavioural effect in a paradigm to be used in the 7T fMRI Experiment 8 – with the hope of thereby maximising any possible associated neural effect in Experiment 8. The present experiment additionally served the purpose of allowing examination, for the first time, of whether the parametric degradation of contingency at test influences the extent of predictive behavioural effects.

It is well established in the animal learning literature that the particular contingency influences learning, such that the mechanisms involved in associating two events depends upon the contingency between event 1 and event 2 (Rescorla & Wagner, 1972). Under this model, expectations can be formed on the basis of learning about the associative strength of two co-occurring events and this therefore determines the predictive power of an event over another. For instance, human instrumental learning studies have shown that learning is dependent on the contingency between actions and outcomes, where learning was observed when the probability of paired presentations was high but reduced as the probability decreased (Cook et al., 2010; Elsner & Hommel, 2004; Shanks & Dickinson, 1991).

It is however less clear how the degradation of contingency influences the degradation of learnt associations. There have been claims that, once associations between sensory events have been established, extinction is rare (Dickinson, 1980; Dickinson & Charnock, 1985). Under this account, behavioural eradication is only observed when new contingencies are introduced that present competing associations, and thereby act in opposition to the first-acquired associations. It is assumed that these competing associations are only established with contingencies above 50% and therefore it is possible that less extreme degradation of a learnt contingency may not degrade behavioural effects. However, it is also plausible that every time an event is presented that violates the learnt association, the original expected event becomes less expected because surprising events are thought to elicit the primary source of (new) learning (Rescorla & Wagner, 1972). For instance, if the contingency between two paired events is degraded to 50% after a period of learning that they are perfectly associated, frequently presented unexpected events will elicit rapid learning about this new relationship (Dickinson, 1980; Rescorla & Wagner, 1972). One possible behavioural

outcome to degraded contingencies could therefore result in reduced RT expectation benefits when expected and unexpected are equally likely. However, learning quickly about new relationships that remain predictive (e.g., when contingencies remain high) could sustain RT benefits when actions' predictive power remains strong.

In sum, although definitions of prediction are often operationalised according to the associative relationship between two events, it has not been explicitly examined how the subsequent degradation of contingencies influences behavioural effects of prediction on perception. Experiment 7 therefore trained participants on a perfectly predictive association between action and outcome and degraded the contingency at the test session 24 hrs later to either 50%, 66.7% or 75%, in separate randomly allocated groups. RTs to expected and unexpected events were measured to establish the extent of expectation benefit in each group. It was anticipated that the contingency at test would most likely modulate the magnitude of effects, such that the strength of RT expectation benefit should be modulated by the contingency at test. In other words, the size of RT benefit will increase as the degraded contingency between action and outcome at test increases across groups. Importantly, this experiment provided a pilot for the paradigm to be used in the laminar fMRI paradigm of Experiment 8, aiming to maximise the predictive effects on perception in the behavioural task and thereby also maximising predictive neural effects in Experiment 8.

5.1.1. Methods

Participants

Twenty-four participants were recruited in each of the low (12 female, mean age = 27.17 years, $SD = 7.03$), medium (15 female, mean age = 29.67 years, $SD = 7.07$) and high (15 female, mean age = 25.92 years, $SD = 5.75$) contingency groups, totalling

seventy-two participants overall. Participants were recruited via Prolific, completed the experiment online, and were paid a small honorarium for participation. One participant was a replacement for a participant who scored worse than chance accuracy on the task. This online experiment was performed with local ethical committee approval.

Procedure

The procedure was similar to the experimental paradigm in Experiment 6, except for the following changes. The experiment was now programmed to run using Gorilla (www.gorilla.sc) for online experiments, and participants took part on either a laptop or desktop computer. Instruction at the beginning of the experiment requested participants to set their screen brightness to the maximum level in an attempt to approximately equalise viewing conditions. For the main task, participants were instructed to hold down the ‘c’ and ‘m’ keys with their respective index and little fingers on their computer keyboard. As soon as a finger lift was executed, the visual stimulus was presented for 500 ms in an annulus surrounding the response question in shorthand (e.g., ‘CW?’ or ‘CCW?’) in the middle of the screen, and replacing the fixation cross (see Fig. 5.1.A). This response text appeared in synchrony with stimulus onset. Participants could respond from the beginning of stimulus onset, enabling reaction time to more tightly couple with perceptual sensitivity, and the question remained on screen until a response was selected – although the stimulus only remained on screen for 500 ms. The stimulus orientation probed by the question was alternated every mini-block of 36 trials. Responses were made using the left thumb on the ‘A’ and ‘Z’ keys on the keyboard for ‘yes’ or ‘no’, respectively. Following response, a variable ITI of 1500-3000 ms was presented before the start of the next trial.

Participants performed this task over two sessions. In the first session, stimulus orientation was perfectly predicted by the action executed. In the second test session 24 hrs later, participants were allocated into one of three groups where their action now predicted the stimulus orientation with either 50%, 66.7% or 75% validity. There were 360 main experimental trials in each session. Participants completed 32 practice trials before proceeding to the main trials in the training session.

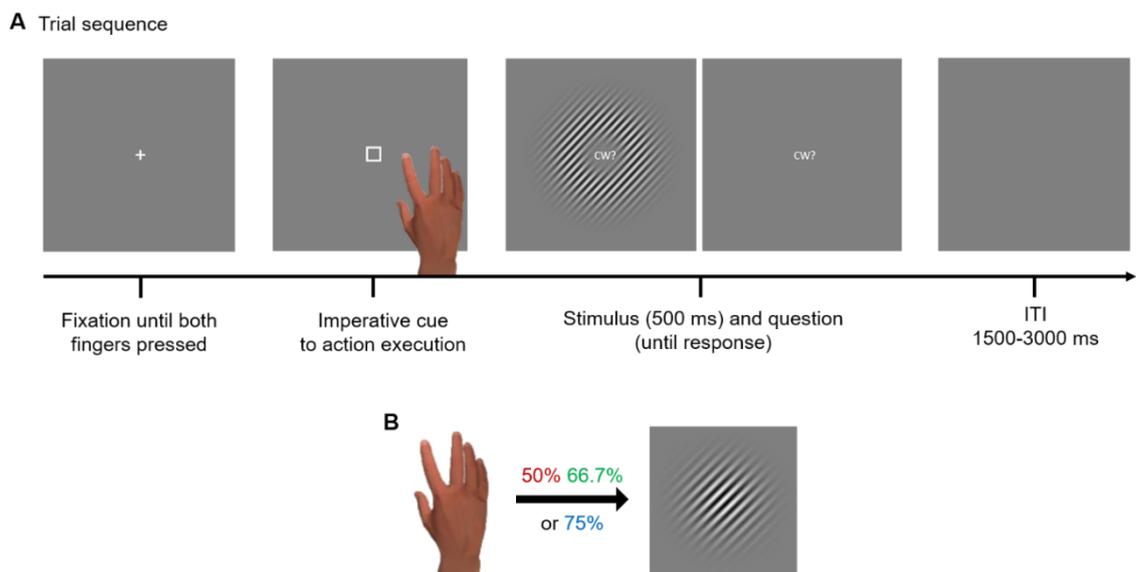


Figure 5.1. Experimental paradigm. (A) A centrally presented visual cue instructed participants to abduct either their index or little finger. In training, each finger lift perfectly predicted an oriented stimulus and participants were required to respond (yes/no) to the question in the middle of the screen – indicating whether the stimulus was clockwise (‘CW?’) or counter-clockwise (‘CCW?’). (B) In the test session, participants performed the same task but were allocated into one of three groups where actions now predicted stimulus orientations with either 50%, 66.7% or 75% validity.

5.1.2. Results and Discussion

To test whether predicted action outcomes resulted in behavioural benefits for expected relative to unexpected events, RT data and proportion of errors in the test session were compared using 2 x 3 mixed ANOVAs with the factors Expectation (expected, unexpected) and Contingency (50, 66.7, and 75). Median RTs and mean proportion of errors were calculated for each individual separately for expected and unexpected events across the three groups. The RT analysis revealed no main effect of Contingency ($F(2, 69) = 1.14, p = .33, \eta_p^2 = .032$), but a significant main effect of Expectation ($F(1, 69) = 31.91, p < .001, \eta_p^2 = .32$), and qualifying this main effect, an interaction between Expectation and Contingency ($F(2, 69) = 9.35, p < .001, \eta_p^2 = .21$). This was reflected by significantly faster responses to expected trials ($M = 734.47$ ms, $SD = 149.42$) than unexpected trials ($M = 780.47$ ms, $SD = 147.21$) in the 75% contingent group ($t(23) = -5.67, p < .001, d = -1.16$), and 66% contingent group (Expected: $M = 685.64$ ms, $SD = 126.61$, Unexpected: $M = 702.57$ ms, $SD = 120.26, t(23) = -3.51, p = .002, d = -.72$) but no difference between expected ($M = 730.15$ ms, $SD = 163.76$) and unexpected ($M = 734.96$ ms, $SD = 169.87$) trials in the 50% contingent group ($t(23) = -.65, p = .522, d = -.13$; see Fig. 5.2.A). Furthermore, a post-hoc 2 x 2 mixed ANOVA was conducted with the factors Expectation and Contingency with only the two levels that presented higher than chance event probabilities (66.7 and 70), to examine whether the significant RT effects in these groups were modulated by the extent of contingency. Indeed this analysis revealed a significant interaction between Expectation and Contingency ($F(1, 46) = 9.49, p = .003, \eta_p^2 = .17$), reflecting larger RT expectation effects in the 75% than the 66.7% contingency group.

A further post-hoc analysis was conducted to examine how effects changed across time. To this end, RT expectation effects (Unexpected RT – Expected RT) were examined

across block in each contingency group in a 2 x 10 mixed ANOVA. This analysis revealed a significant main effect of Block ($F(3.37, 232.78) = 4.13, p = .005, \eta_p^2 = .06$, Greenhouse-Geisser corrected, $\varepsilon = .38$), and Contingency ($F(2, 69) = 8.67, p < .001, \eta_p^2 = .20$), but no significant interaction between Block and Contingency ($F(6.75, 232.78) = 1.67, p = .120, \eta_p^2 = .05$, Greenhouse-Geisser corrected, see Fig. 5.2.D). However, a further analysis of the polynomial trends across block revealed a significant linear interaction between Block and Contingency ($F(2, 69) = 3.16, p = .049, \eta_p^2 = .08$), suggesting that the linear slope across blocks differed depending upon the contingency group. Fig. 5.2.D demonstrates that this likely reflects the steeper decline for RT effects across time in the 75% relative to the other two groups.

A similar pattern was observed in proportion of errors, demonstrating a main effect of Expectation ($F(1, 69) = 32.97, p < .001, \eta_p^2 = .32$), a main effect of Contingency ($F(2, 69) = 5.18, p = .008, \eta_p^2 = .13$), and an interaction between Expectation and Contingency ($F(2, 69) = 8.61, p < .001, \eta_p^2 = .20$). This interaction is generated by a significantly smaller proportion of errors on expected trials ($M = .054, SD = .044$) than unexpected trials ($M = .101, SD = .079$) in the 75% contingent group ($t(23) = 4.22, p < .001, d = -.86$), and on expected ($M = .032, SD = .027$) than unexpected trials ($M = .057, SD = .035$) in the 66% contingent group ($t(23) = 5.35, p < .001, d = -1.09$), but no difference between expected ($M = .043, SD = .029$) and unexpected ($M = .046, SD = .033$) in the 50% contingent group ($t(23) = .59, p = .560, d = -.12$; see Fig 5.2.B). A post-hoc 2 x 2 mixed ANOVA was conducted on the accuracy data to determine whether there was an interaction between Expectation and Contingency for only the two higher contingency groups, but this analysis revealed that there was not ($F(1, 46) = 3.34, p = .074, \eta_p^2 = .07$).

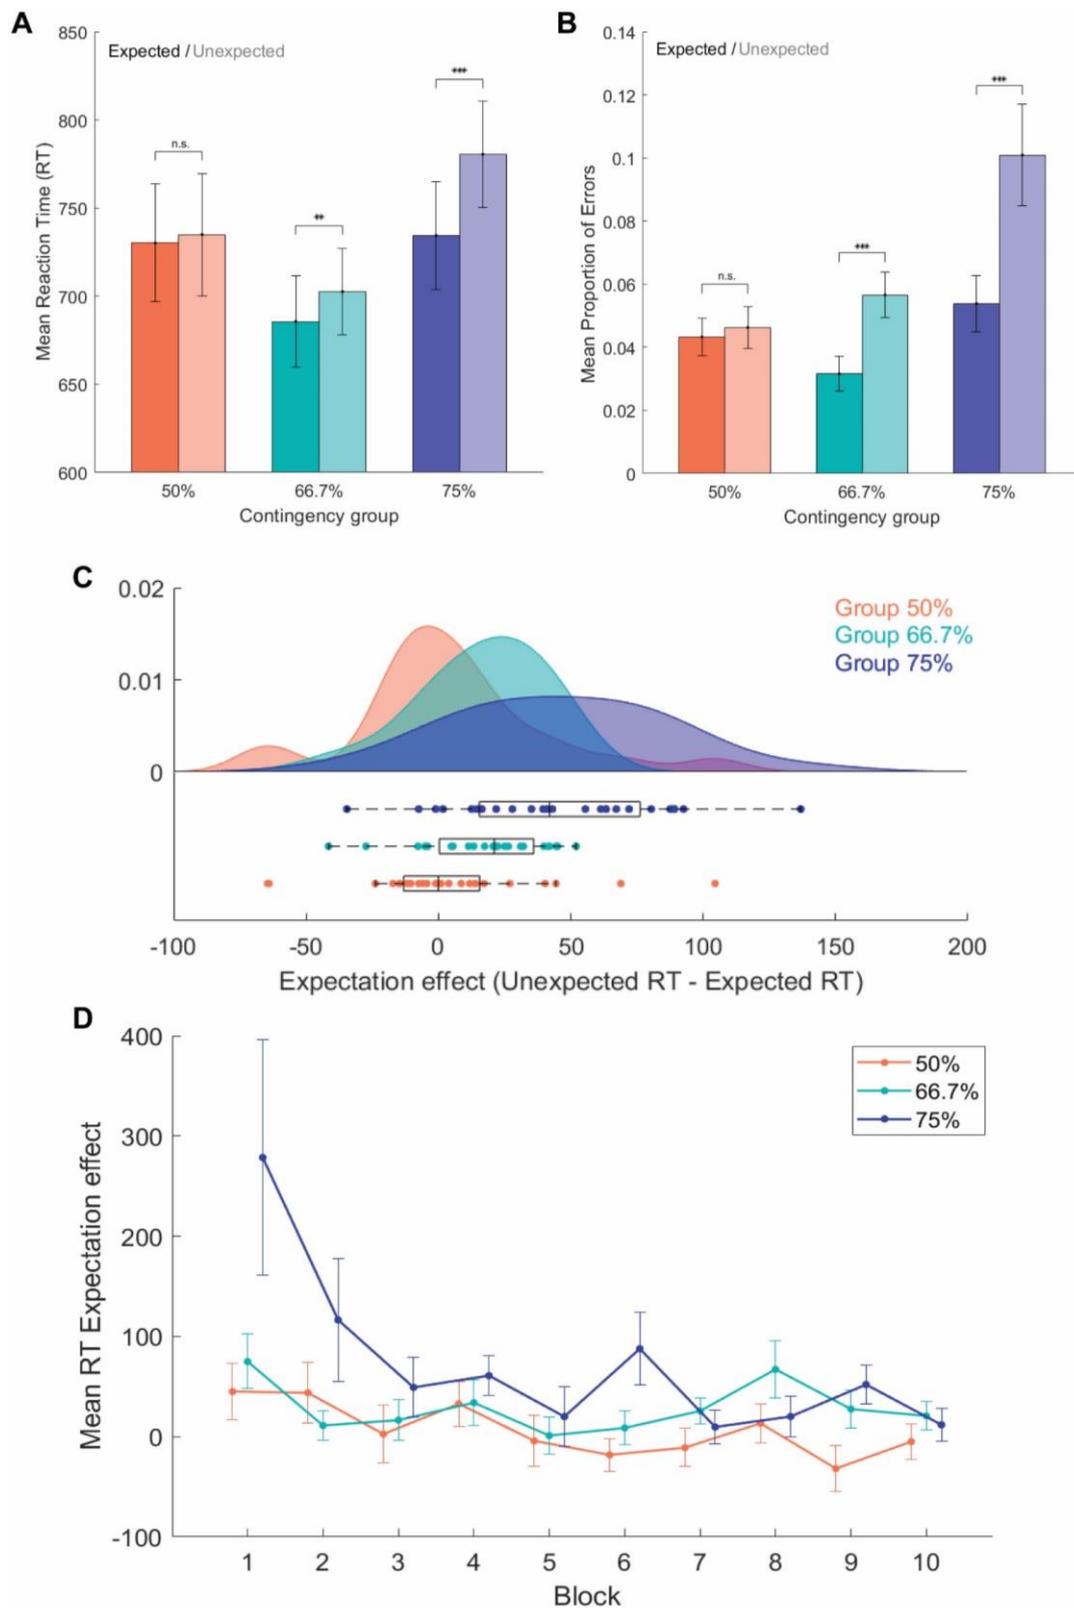


Figure 5.2. (A) Mean RTs (\pm SEM) were faster for expected than unexpected trials in both 66% and 75% contingent groups but not when 50% contingent. (B) Mean proportion of errors (\pm SEM) were lower for expected than unexpected trials in both

66% and 75% contingent groups but not when 50% contingent (** $p < .01$, *** $p < .001$). (C) RT expectation effects (Unexpected RT – Expected RT) were calculated for the 50% [peach], 66% [teal], and 75% [blue] contingent data and are plotted with raincloud plots (Allen et al., 2019) displaying probability density estimates (upper) and box and scatter plots (lower). Boxes denote lower, middle and upper quartiles, whiskers denote 1.5 interquartile range, and dots denote difference scores for each participant ($N = 24$ per group). Positive effects indicate faster responses to expected events relative to unexpected events. (D) Mean RT effects (Unexpected RT – Expected RT, \pm SEM) by mini-block across the test session.

These results show that participants are faster, and less error prone, when responding to expected events that had a high contingency between action and outcome, and such effects were directly modulated by the degradation of contingency. Specifically, RT benefits were strongest when the contingency remained the highest at test (75%) and became weaker as the contingency between action and outcomes was degraded.

Interestingly, this pattern changed over time such that RT effects were numerically stronger for all groups at the beginning of the test block, and could perhaps reflect quick learning of the new relationship between action and stimulus orientation in subsequent blocks.

Interestingly, the modulation of RT effects by degraded contingency across blocks at test could be supported by animal learning theories that propose that learning is facilitated by the surprising nature of an event that violates learned associations (Rescorla & Wagner, 1972). For example, in the low contingency group early presentation of unexpected events may have induced fast learning that action now predicts the previous expected outcome at chance, and therefore reduced RT benefits for ‘expected’ events as early as the first block. Moreover, large initial RT benefits in the high contingency group may be explained by stronger reliance on the previously learned

perfectly predictive relationship, because learning of the new relationship is slower when unexpected events are presented less frequently and actions still hold high predictive power. Although unexpected events in this group may elicit the highest surprise and therefore should facilitate new learning (Press et al., 2020b; Rescorla & Wagner, 1972), such learning perhaps does not occur as rapidly in the first block because fewer of these events will have been presented relative to in the other groups. Any RT effects that remain after the initial drop may thereby reflect benefits that are proportionate in magnitude to the remaining predictive power of action upon new learning – and would support the finding that RT effects were generally greater in the 75% group across the experiment.

This experiment therefore suggests that the degradation of contingencies at test modulates predictive influences on perception, such that effects are stronger when events are associated with higher contingencies and weakest when they become non-contingent with action. This is consistent with theoretical accounts that suggest that predictions that are formed on the basis of statistical learning influence perception, and suggests that stronger predictive relationships will produce the largest perceptual influences. The aim of this experiment was primarily to examine the contingency level that would provide the largest behavioural effects, on the premise that this might also manifest in greater predictive neural effects, and so the 75% contingency group was selected to measure the laminar profile of expected and unexpected events in Experiment 8.

5.2. Experiment 8

In the last decade, studies investigating prediction mechanisms across cortical layers in humans have used ultra-high resolution 7T MRI to examine laminar activation patterns in V1. In general, the aims of these studies have been to examine evidence for the idea that P and PE units operate in different layers of the cortex, supporting hierarchical predictive coding theories. For instance, these theories suggest that predictions or ‘hypotheses’ about incoming input are conveyed through descending feedback connections originating in the deep layers of higher-order regions and are passed to lower-order agranular (superficial and deep) layers. On the other hand, error signals are thought to be conveyed through ascending feedforward connections that target the middle layers of lower-order regions but are passed up to superficial layers to process the mismatch between concurrent input and predictions e.g., the prediction error (Bastos et al., 2012; Friston, 2005; see 5.3.A).

As discussed, studies have provided evidence for the idea that predictions or hypotheses about our incoming sensory input are represented in the deep layers of V1, in the absence of bottom-up input. For example, illusory contours of visual Kanizsa figures (Kok, Bains, Van Mourik, et al., 2016; Lee & Nguyen, 2001) and events that are expected but subsequently omitted (Aitken et al., 2020) have been shown to be represented in the deep layers of V1. Another study examined contextual feedback processes in the absence of bottom-up feedforward input by presenting visual scenes in which the lower quadrants of images were masked so that a sub-region of the visual field was occluded (Muckli et al., 2015). Muckli and colleagues (2015) used MVPA techniques to decode the content of visual scenes from patterns of activity in V1, and found superior classification of the content of these occluded scenes only in the most superficial cortical layers. They suggest that these effects reflect contextual feedback

information about the image scene, but this could be in contrast to similar contextual ‘filling-in’ feedback effects that are found in the deep layers of V1 (Kok, Bains, Van Mourik, et al., 2016; Lee & Nguyen, 2001).

However, it is quite plausible that a difference in methods generated these differences in findings. The studies by Kok et al. (2016) and Aitken et al. (2020) used a spatial regression approach to de-correlate responses across layers and demonstrated increased BOLD responses in deep layers that were not also present in middle or superficial layers. On the other hand, Muckli et al. (2015) adopted multivariate pattern techniques to examine the extent to which occluded content could be decoded from patterns of visual activity. This difference is particularly notable because studies have revealed that decoding accuracy can increase towards superficial cortical layers (Moerel et al., 2017), caused by an increase in signal from venous draining towards the pial surface (Koopmans et al., 2010; Lawrence, Formisano, et al., 2019). Specifically, gradient echo BOLD signal is known to exhibit strong contributions from large veins situated perpendicular to the cortical surface, as venous blood is drained from lower to upper cortical layers (Duvernoy et al., 1981; Markuerkiaga et al., 2016; Stephan et al., 2019). This means that the BOLD response at deeper cortical layers will contribute to responses in more superficial layers, making interpretations about layer-specific BOLD responses in superficial layers complex (Kok, Bains, Van Mourik, et al., 2016). Nevertheless, both findings would remain consistent with the functional architecture of feedback projections, known to have dense terminations in deep and superficial layers while avoiding middle layers (Anderson & Martin, 2009; Rockland & Virga, 1989; van Kerkoerle et al., 2017).

Studies aiming to mitigate such contributions of draining towards superficial layers have made comparisons between stimulus-specific BOLD responses within layers, thereby examining the relative differences between signals when draining influences should be equivalent (Aitken et al., 2020; Albers et al., 2018; Lawrence et al., 2018; Lawrence, Norris, et al., 2019). Specifically, this technique involves subtracting activity in voxels not preferring the presented orientation (e.g., tuned to CCW orientations when CW is presented) from activity in voxels that prefer the presented orientation (e.g., tuned to CW orientations) – obtaining a measure that represents an orientation-specific signal, instead of absolute signal that is more susceptible to draining effects. Notably, several laminar fMRI studies examining evidence for the presence of distinct top-down processes have opted to examine activity patterns in this way, in particular when events are expected in the absence of bottom-up input. This choice is unsurprising considering bottom-up input leads to signal propagation across the cortical column, reflecting message passing between the layers when both top-down and bottom-up processes occur at the same time (Lawrence, Norris, et al., 2019; Self et al., 2013). Consequently, measuring signals with methods that have poor temporal resolution such as fMRI increases the likelihood that such propagation may mask individual top-down influences across the column that cannot be temporally distinguished.

Regardless of these methodological challenges, it is essential that attempts are made to overcome them if we are to determine the possibility of PE encoding in superficial layers, because error encoding cannot be assessed without presenting bottom-up input that violates predictions. Despite the fact that no studies have yet examined this question, some recent studies have attempted to provide evidence for bottom-up processing in middle layers that is distinct from top-down processing. For example, a recent study measuring orientation-specific BOLD signals cleverly designed an

experiment to examine top-down influences of feature-based attention orthogonally to bottom-up contrast modulations (Lawrence, Norris, et al., 2019). Participants viewed stimuli that comprised of bars of CW or CCW orientation and were required to respond to any changes in bar width in the cued orientation, while on each trial stimulus contrast was also modulated to be high (80%) or low (30%). The same spatial regression approach (Kok, Bains, Van Mourik, et al., 2016; van Mourik, van der Eerden, et al., 2019) was adopted to de-correlate responses in each layer and reduce influence of venous draining towards the pial surface. Distinct activity patterns were observed for top-down attention relative to bottom-up contrast manipulations, namely, attended oriented bars elicited larger orientation-specific BOLD responses across all layers relative to unattended stimuli, but this effect was strongest in the superficial layers. On the other hand, modulations of stimulus contrast were evident most strongly in the middle layers. These findings are perhaps consistent with findings that bottom-up input primarily targets the middle layers of lower-level regions (Hubel & Wiesel, 1972), while top-down processes target agranular layers (Felleman & Van Essen, 1991), but the specific layer targeted may depend on the type of top-down manipulation (van Kerkoerle et al., 2017).

Evidence for *predictive* top-down influences in the presence of presented stimuli are more scarce, and the recent study by Aitken and colleagues (2020) revealed that the orientation-specific BOLD signal of a presented expected stimulus was equivalent across all layer bins – in contrast to deep layer activation of expected but omitted stimuli – likely reflecting activation that propagated across all layers. Self et al. (2013) assessed layer-specific neural responses in macaques during presentation of a figure-ground segregation task in the presence of bottom-up stimulation and found evidence for figure filling-in in layer 5 (deep), as well as superficial layers I and the upper part of

layer II. Kok et al. (2016) examined activity patterns elicited by Kanizsa triangles in stimulus regions that produced illusory contours as well as on Pac-Man shapes that induced the illusion but provided bottom-up input. The latter were found to induce suppression of activity across all layers relative to identical bottom-up input that did not contextually induce illusory figures, but this effect was largest in the superficial layers. However, activity patterns for top-down processes in illusory contour regions evoked increased activity in the deep layers relative to regions that contained no illusory contour – together somewhat consistently with Self et al.’s (2013) findings in the presence of bottom-up input. These comparisons reveal distinct laminar profiles in superficial and deep layers for feedback processes that appear dependent upon whether bottom-up inputs interact with top-down processes (Kok, Bains, van Mourik, et al., 2016). It is possible that such differences when stimuli are, or are not, presented across layers provide preliminary support for the hypothesis that distinct prediction-related processes operate in superficial and deep layers, considering that bottom-up inputs are thought to convey PE signals that are matched with hypotheses in the superficial layers (Bastos et al., 2012; Friston, 2008).

However, due to the lack of direct evidence indicating that these bottom-up processes reflect PE processing when an event is presented and thus violates predictions, such a hypothesis remains unsupported. Most importantly, there have been to date no laminar fMRI studies that present unexpected stimuli – notably likely due to the issues discussed surrounding propagation of signals throughout the column. One way to reduce such issues could be to compare the relative activation across layers when expected and unexpected events are always presented as bottom-up input. Experiment 8 therefore investigates whether P and PE populations operate with distinct functional roles and across different cortical layers using 7T fMRI – exploring the primary question

surrounding whether such distinct processes could account for the discrepancies across neural findings. This question will be investigated in early visual region V1 and additionally V2 which has also been previously shown to encode orientation (Anzai et al., 2007; Lawrence, Norris, et al., 2019). Further, due to the close spatial proximity of these areas, both could be captured within the tight field of view.

Participants were trained with a perfect contingency between actions and stimulus orientation in similar ways to described in Experiment 7. Twenty-four hours later at test (during scanning) the contingency was degraded to 75% to measure neural responses to expected and unexpected events. Orientation-specific activity patterns were measured across deep, middle, and superficial cortical layers, and the same spatial regression approach was adopted to de-correlate laminar signals across layers as has been used in recent laminar studies (Aitken et al., 2020; Albers et al., 2018; Lawrence et al., 2018; Lawrence, Norris, et al., 2019; van Mourik, van der Eerden, et al., 2019). It was hypothesised that expected events will evoke orientation-specific BOLD signals across all layers in a similar manner to those observed in Aitken et al. (2020), considering that the only large difference between paradigms is that actions (rather than coloured fixation dots) now predicted sensory outcomes. However, if PE signals are indeed represented in superficial layer populations, as the central tenet of predictive coding theories would propose, then unexpected events should evoke larger orientation-specific signals in superficial layers relative to expected events. PE orientation-specific activity should therefore be smallest in deep layers and lower than activity elicited by expected events – this would reflect greater top-down activation of P populations in deep layers, and such a pattern would equate to effects of expectation suppression and enhancement in superficial and deep layers, respectively (see Fig. 5.3.B).

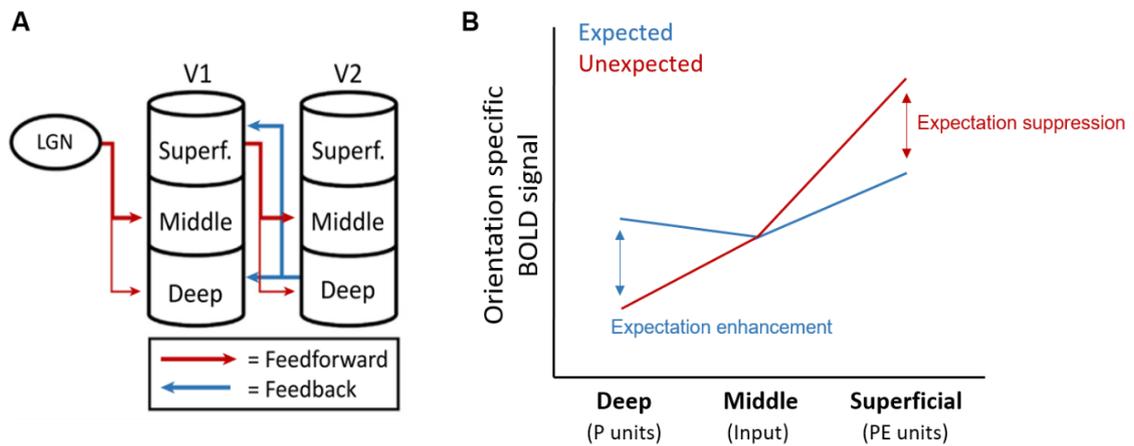


Figure 5.3. (A) Schematic representation of proposed extrinsic feedforward (red) and feedback (blue) connections across layers between the LGN, V1 and V2; taken from Lawrence, Formisano, et al. (2019). (B) Hypothesised orientation-specific BOLD (preferred – non-preferred) responses to expected and unexpected events across layers, and according to key assumptions of predictive coding theories. Arrows highlight a possible reconciliation to the paradox that expectation suppression and enhancement (or ‘sharpening’) are represented in different cortical layers.

5.2.1. Methods and Analyses

Participants

Fifteen participants were recruited (12 female, mean age = 26.60 years, $SD = 3.22$) from Birkbeck College and UCL and paid a small honorarium for their participation. All participants reported normal or corrected to normal vision and had no history of psychiatric or neurological illness. One participant’s data was excluded due to a technical error during acquisition, which meant that event onsets in one run could not appropriately be modelled. The experiment was approved by the UCL ethics committee. This sample size was limited by COVID-19 restrictions and therefore does not fulfil power calculations for medium effect sizes – where the intended sample size was 20 participants to match the sample collected in Experiment 6 and previous studies (Aitken et al., 2020; Kok, Jehee, et al., 2012).

Procedure

The scanning procedure was almost identical to the experimental paradigm described for expected and unexpected events in Experiment 6, except for the following changes. During scanning stimuli were presented on an LCD monitor viewed through a mirror at a viewing distance of 91 cm. Actions predicted expected events with 75% validity in order to maximise influences of prediction during scanning. Response to the question screen were now required within 1000 ms. Participants completed the main task in four scanning runs that contained 96 trials each, and a 30 s break was presented mid-way through each run.

The training procedure adopted the same paradigm as the online Experiment 7, and participants took part on either a laptop or desktop computer 24 hrs before the scanning session. The only difference with respect to Experiment 7 was that responses were measured after a variable ISI of 300-500 ms, rather than from stimulus onset. This protocol is equivalent to Experiment 6. A short refresher of the training session was also presented immediately before the scanning session, which had the same set-up but participants now executed actions using MR compatible button boxes outside of the scanner.

There were 384 main experimental trials in the scanning session, 360 online training trials and 192 refresher training trials. Participants completed 32 practice trials before proceeding to the main trials in the initial training session.

During scanning, participants completed a functional localiser task in an additional scanning run at the end of the main experimental task. This task presented flickering grating stimuli at approximately 1.8 Hz that were identical to the gratings presented in the main experiment except that they were now presented at 100% contrast, and in

blocks of approximately 14 seconds each. Each block containing flickering gratings was preceded by a blank screen containing only the fixation cross for the same duration. In each stimulus block gratings were fixed to either CW or CCW orientations, and this order was pseudorandomised. The task required participants to respond by pressing any button when the central fixation cross changed colour from white to a light grey, ensuring that fixation remained central. In total, 16 blocks of flickering grating stimuli were presented, eight of each orientation. This procedure was almost identical to the functional localiser presented in (Aitken et al., 2020).

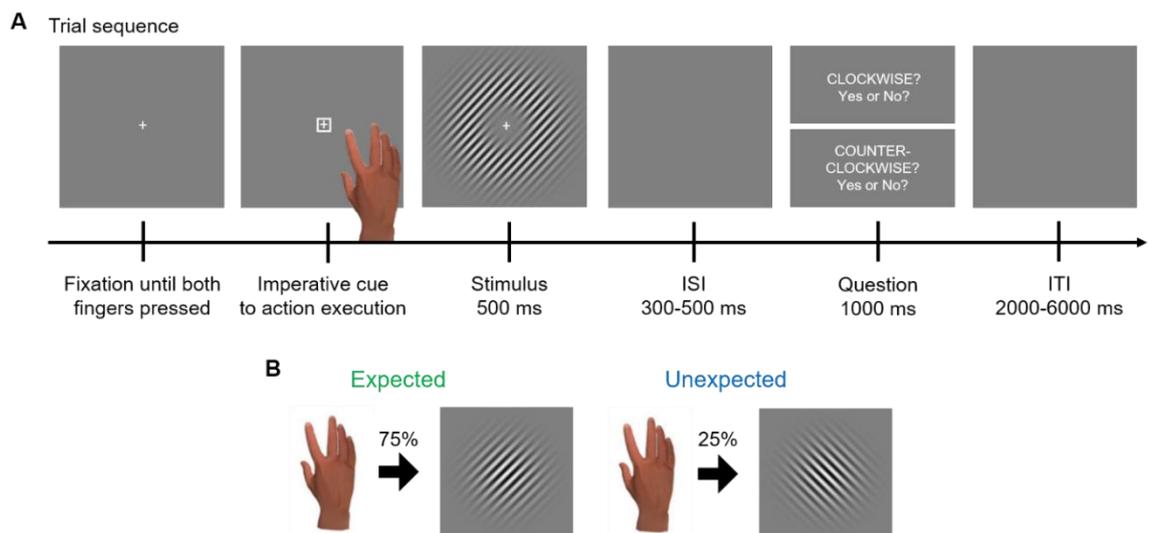


Figure 5.4. Experimental paradigm. (A) The same trial sequence as Experiment 6. A centrally presented visual cue instructed participants to abduct either their index or little finger. Each finger abduction predicted an oriented stimulus and participants were required to respond (yes/no) to whether the stimulus was clockwise or counter-clockwise. In training, actions perfectly predicted the stimulus orientation that was presented. (B) In the scanning session 24 hrs later, participants completed the same task but now actions predicted the oriented stimulus with 75% validity.

Behavioural analyses

RT data was collected for responses to expected and unexpected stimuli in the test session and median RTs were calculated across all trials separately for each condition, for each participant. Similarly, the proportion of correct responses was analysed for expected and unexpected conditions for each participant.

Image Acquisition

Images were acquired using a 7T Magnetom MRI scanner (Siemens Healthcare GmbH, Erlangen, Germany) using a 32-channel head coil at the WCHN, UCL. Functional images were acquired using T2*-weighted 3D gradient-echo EPI sequence (3,552 ms volume acquisition time, TR = 74 ms, TE = 26.95 ms, 48 slices, 15° flip angle, voxel size: 0.8 x 0.8 x 0.8 mm, field of view: 192 x 192 x 39 mm). Structural images were acquired using a Magnetization Prepared Two Rapid Acquisition Gradient Echo (MP2RAGE) sequence (TR = 5,000 ms, TE = 2.60 ms, TI = 900 ms, 240 slices, voxel size 0.7 × 0.7 × 0.7 mm, 5° flip angle, field of view 208 × 208 × 156 mm). The first two volumes acquired in every run were opposite phase-encoded (e.g., one ascending and one descending) to measure the distortions present in each run. These volumes can be used for blip-up blip-down distortion correction analyses but such analyses are not reported here. A high-contrast whole-brain EPI scan was acquired for the possibility that co-registration might be improved by using it as an intermediate for registering the high-resolution anatomical with the low-resolution functional images, but again this type of co-registration is not reported in this thesis.

fMRI data preprocessing and coregistration of cortical surfaces

Preprocessing of the images was conducted in SPM12 and Freesurfer

(<http://surfer.nmr.mgh.harvard.edu/>). Functional images were cropped to select only the occipital lobe, to account for distortions present in the frontal lobes. These cropped functional images were then spatially realigned to the mean image within runs, but also across runs, correcting for head motion. The temporal signal-to-noise ratio (tSNR, defined as mean signal/*SD* over time) was calculated before and after spatial realignment and was found to be significantly higher after (14.25 ± 1.05 , mean \pm *SD* over participants) than before (10.36 ± 1.31) realignment ($t(13) = -16.89$, $p < .001$).

The realigned functional images were then coregistered to the cortical surfaces estimated in participants' MP2RAGE scans in several steps. This procedure was identical to that described in Aitken et al. (2020). First, boundaries between grey matter (GM), white matter (WM) and cerebrospinal fluid (CSF) were detected and estimated using Freesurfer on reconstructed structural scans (skull removed). Boundaries were manually corrected to remove any dura that was inaccurately classified as part of the GM surface. A rigid body boundary-based registration (BBR; Greve & Fischl, 2009) was used to register GM boundaries to the mean functional image, and a further recursive boundary-based registration (RBR; van Mourik, Koopmans, et al., 2019) applied the BBR recursively to portions of cortical mesh in 6 iterations (see Fig. 5.5.A for an example boundary-based coregistration).

ROI definition

The ROI definition technique used was identical to previous studies that determined orientation-specific BOLD activity across layers (Aitken et al., 2020; Albers et al.,

2018; Lawrence et al., 2018; Lawrence, Norris, et al., 2019). Measuring orientation specific activation allows us to control for limitations of laminar fMRI that signals increase towards superficial layers due to venous blood draining towards the surface. This method should allow examination of orientation-specific activity changes across layers that should contain equivalent contributions of draining. Freesurfer was used to define the ROIs, V1 and V2, based on anatomical landmarks in the MP2RAGE scans, similarly to in Experiment 6. However this time, given that the protocol included a functional localiser, ROIs were restricted to voxels from the preprocessed functional localiser data that were most active during blocked presentation of the stimuli. This was achieved by modelling regressors for blocks of CW and CCW stimuli against baseline in a temporal GLM to identify voxels that expressed a significant response these stimuli ($t > 2.3$, $p < 0.05$; 4884.67 \pm 1845.74 voxels, mean \pm SD over participants). The orientation preference of each of these voxels was then calculated by contrasting CW and CCW regressors to select the 500 voxels that were most CW preferring ($t > 0$) and CCW preferring ($t < 0$). Separate masks were created for these CW and CCW preferring voxels in V1 and V2. The timecourse of each voxel in these masks was then z-scored and multiplied with its t value from the orientation contrast, to weight the results towards voxels with the largest orientation preferences (Aitken et al., 2020; Lawrence et al., 2018; Lawrence, Norris, et al., 2019).

Cortical layers definition & extraction of layer-specific timecourses

The level set method was used to divide GM into three equivolume layers, adopting a method identical to Aitken et al. (2020) for cortical layer definition and is outlined in detail in van Mourik, van der Eerden, et al. (2019). This method was used to separate out five cortical bins (3 GM, WM and CSF) and determine three GM layers (deep,

middle, and superficial) by calculating two intermediate surfaces between the WM and pial boundaries, and provided the foundation for the following laminar analyses. In human V1, these three layer bins have been suggested to correspond to histological layers 1 to 3, layer 4, and layers 5 and 6, respectively (Aitken et al., 2020).

To extract signals from different layers, a spatial GLM approach was adopted to determine the proportion of each voxel's overlap to each of the three grey matter bins in each ROI mask. This approach has previously been successfully used to de-correlate laminar activity profiles in visual areas in tasks of working memory, attention and prediction (Aitken et al., 2020; Kok, Bains, van Mourik, et al., 2016; Lawrence et al., 2018; Lawrence, Norris, et al., 2019; van Mourik, van der Eerden, et al., 2019), and the present analysis follows an identical method. Described in short, a laminar matrix is designed that represents the distributions of voxels across layers in each ROI, and can therefore be used in a spatial GLM to separate out BOLD signals for each voxel into the specified layer bins. The GLM of spatially distributed responses unmixes signals from the different layers (Kok, Bains, van Mourik, et al., 2016), and produces timecourses for each layer for each ROI in each functional run.

Estimation of laminar responses

Timecourses from the three GM layers of interest were then selected for a temporal GLM in SPM12 to estimate laminar responses for each of the main experiment conditions. Regressors of interest were modelled for each experimental condition (expected, unexpected) and each orientation that was presented (CW, CCW), resulting in four regressors of interest for each scanning run. These regressors were modelled by convolving stick functions representing the onset of presented stimuli in the main

experiment with the canonical haemodynamic response function. Nuisance regressors modelled head motion parameters and their temporal derivatives.

To investigate how patterns of activity differed across layers, orientation-specific BOLD responses were calculated for each ROI (e.g., mask for voxels preferring CW in V1), in each run, by subtracting the estimated BOLD activity from conditions that presented the non-preferred orientation (e.g., an expected CCW stimulus) from activity when the preferred orientation was presented (e.g., an expected CW stimulus). This was done correspondingly for expected and unexpected conditions, and responses were averaged across CW and CCW masks resulting in orientation-specific BOLD responses (preferred activity – non-preferred activity) for each of the conditions across each layer.

These activity patterns were then analysed with 2 x 3 repeated measures ANOVAs with the factors experimental condition (expected, unexpected) and cortical layer (deep, middle, superficial), separately for V1 and V2. The main hypothesis that distinct effects of expectation will be present in deep and superficial layers was tested by the interaction between Expectation and Cortical layer. Follow up paired samples t-tests examined differences between expected and unexpected activity in each layer for each ROI.

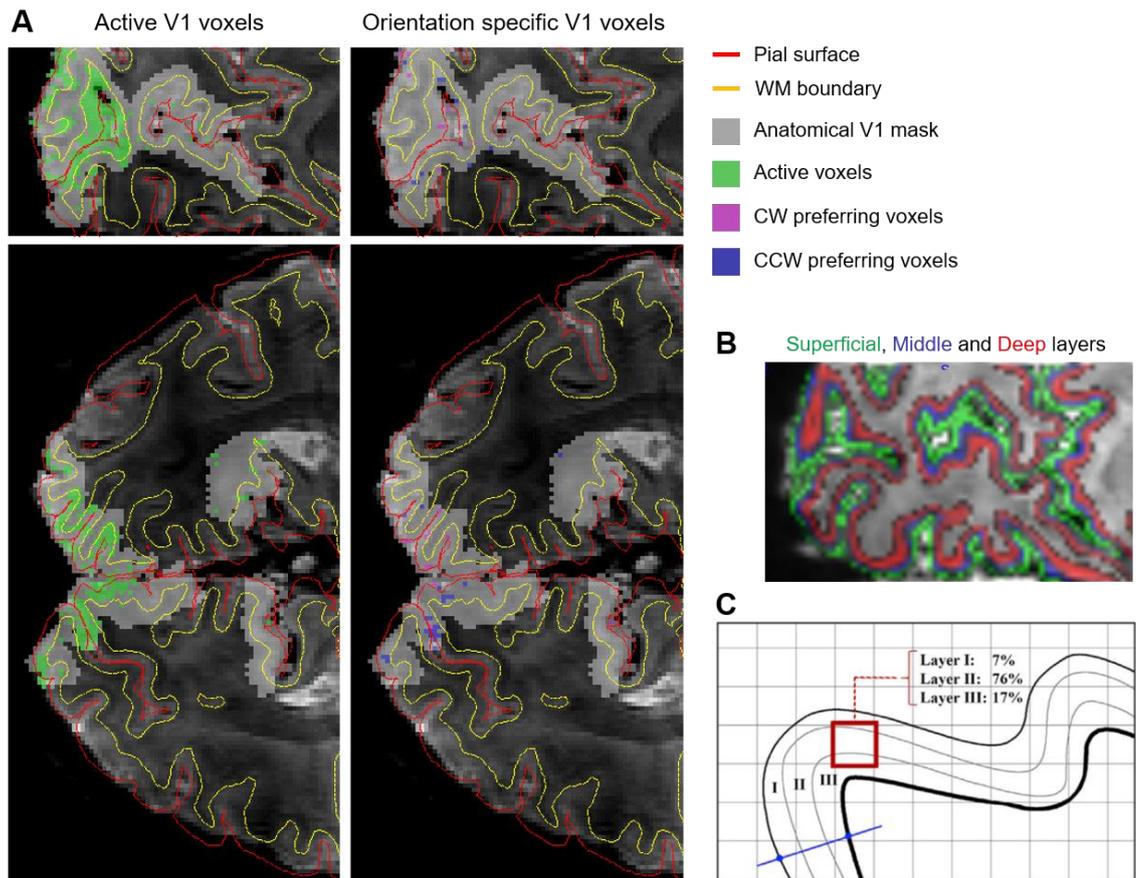


Figure 5.5. (A) Visualisation of the selected anatomical V1 ROI (light grey) on a mean functional image of an example participant. Overlaid red and yellow lines represent co-registered anatomical WM (yellow) and pial surface (red) boundaries to the mean functional image. The left panel shows voxels that were significantly active to the presented stimuli in the functional localiser task (green). The right panel shows the voxels within these active voxels that were strongly tuned to CW (pink) or CCW (blue) orientations. (B) A mean functional image overlaid with distributions of voxels in superficial (green), middle (blue) and deep (red) layers of the GM. (C) A schematic representing the spatial regression approach used to determine the volume distribution of a selected voxel (e.g., red square) over superficial (I), middle (II), and deep (III) cortical layers to estimate laminar time-courses. Taken from Aitken et al. (2020).

5.2.2. Results and Discussion

Behavioural analysis

RT analyses revealed no significant difference between response times for expected ($M = 604.64$ ms, $SD = 66.14$) and unexpected ($M = 611.68$ ms, $SD = 62.71$) trials ($t(13) = -1.05$, $p = .314$, $d = -.28$). Participants demonstrated a trend to be more accurate on expected ($M = .97$, $SD = .03$) trials than unexpected ($M = .96$, $SD = .04$) trials ($t(13) = 1.87$, $p = .084$, $d = .50$).

Laminar responses of Expectation

Analysis of orientation-specific BOLD responses to expected and unexpected events across cortical layers revealed no significant interaction between Expectation and Cortical layer in V1 ($F(2,26) = .013$, $p = .987$, $\eta_p^2 = .001$) or V2 ($F(2,26) = .042$, $p = .959$, $\eta_p^2 = .003$), which were the primary effects of interest. There were also no main effects of Expectation (V1: $F(1,13) = .14$, $p = .712$, $\eta_p^2 = .01$, V2: $F(1,13) = .41$, $p = .534$, $\eta_p^2 = .03$) or Cortical layer (V1: $F(1.35,17.5) = .49$, $p = .548$, $\eta_p^2 = .04$, Greenhouse-Geisser corrected, $\varepsilon = .67$, V2: $F(1,13) = .026$, $p = .975$, $\eta_p^2 = .002$).

Although responses to expected events were numerically larger than unexpected events in deep (V1, Expected: $M = -.04$, $SD = .77$, Unexpected: $M = -.28$, $SD = 3.25$; V2, Expected: $M = .72$, $SD = 1.57$, Unexpected: $M = .29$, $SD = 2.73$), middle (V1, Expected: $M = .33$, $SD = 1.16$, Unexpected: $M = -.004$, $SD = 2.18$; V2, Expected: $M = .90$, $SD = 1.78$, Unexpected: $M = .30$, $SD = 2.30$) and superficial (V1, Expected: $M = .10$, $SD = 1.55$, Unexpected: $M = -.28$, $SD = 3.25$; V2, Expected: $M = .82$, $SD = 1.99$, Unexpected: $M = .26$, $SD = 2.81$) layers (see Fig. 5.6.), these differences were not significant at any layer in either ROI (V1, deep: $t(13) = .26$, $p = .801$, $d = .07$, middle: $t(13) = .50$, $p = .625$, $d = .183$, superficial: $t(13) = .35$, $p = .734$, $d = .09$; V2, deep: $t(13) = .46$, $p = .656$,

$d = .12$, middle: $t(13) = .841$, $p = .416$, $d = .23$, superficial: $t(13) = .54$, $p = .596$, $d = .15$).

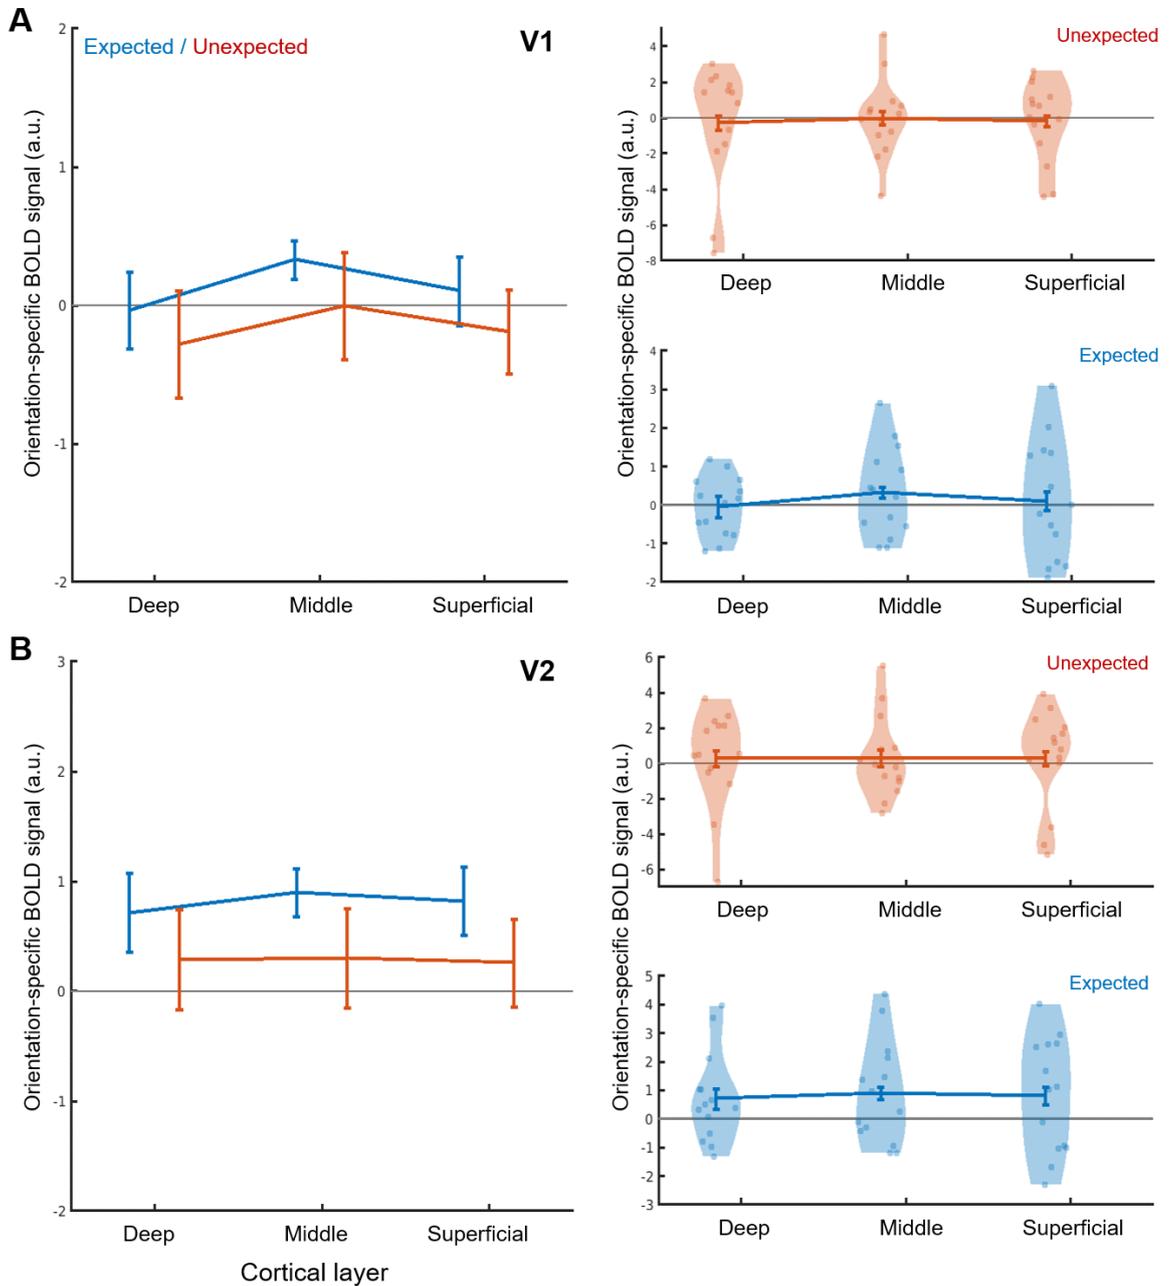


Figure 5.6. (A) Orientation-specific BOLD responses in V1 across cortical layers (deep, middle, and superficial) for expected (blue) and unexpected (orange) events. The same effects are shown averaged over all participants (left panel) as well as presenting individual participants' data points (right panel). (B) Presents the same as (A) but for orientation-specific BOLD responses in V2. Error bars indicate within-subject SEM.

These results may suggest that processing of expected and unexpected events does not differ across layers, likely reflecting signal propagation across all layers when top-down and bottom-up processes interact. However, unfortunately any conclusions at this stage must be highly tentative for two main reasons. Firstly, this study is underpowered and the remaining participants will be tested once COVID-19 restrictions are lifted.

Secondly, this was one of the first two studies to be conducted (in parallel) on the new 7T MRI scanner at the WCHN (UCL) and it is possible that there remains some sub-optimality in the acquisition sequences and analysis pipeline at present, that are likely to be improved following interactions with the physicists. However, this means that the layer-specific orientation signals are difficult to interpret at this stage (see discussion of some of these issues below).

A further analysis was run to examine the univariate BOLD signal across layers before orientation preference subtractions were computed, to determine whether patterns were consistent with the known changes in signal across layers. BOLD responses across layers were examined in three-way repeated measures ANOVAs with the factors Expectation, Preference (preferring, non-preferring) and Layer, separately for each ROI. Importantly, a main effect of Layer was hypothesised to reflect an increase in BOLD signal from deep to superficial layers induced by venous blood draining to the pial surface (Koopmans et al., 2010; Markuerkiaga et al., 2016). This is indeed what was observed in both V1 ($F(2, 26) = 6.51, p = .005, \eta_p^2 = .33$) and V2 ($F(2, 26) = 4.25, p = .025, \eta_p^2 = .25$, see Fig. 5.7.). Bonferroni-corrected pairwise comparisons revealed that these main effects reflected significantly lower BOLD signal in deep (V1: $M = .87, SD = 4.25$, V2: $M = 1.50, SD = 5.61$) than superficial (V1: $M = 4.42, SD = 7.09$, V2: $M = 4.78, SD = 7.24$) layers in both V1 ($p = .011$) and V2 ($p = .030$), and in deep than middle (V1: $M = 3.07, SD = 3.92$, V2: $M = 3.32, SD = 4.70$) layers in V1 ($p = .049$) but

not V2 ($p = .185$), and no difference between superficial and middle in both V1 ($p = .773$) and V2 ($p = .898$), regardless of expectation or preference. No additional significant main or interaction effects were observed and are therefore not reported here for simplicity and because they were not relevant to the question of interest. These results therefore demonstrate laminar-specific BOLD patterns in V1 and V2 that are consistent with patterns that would be predicted across layers in these regions according to venous blood draining towards the cortical surface (Koopmans et al., 2010; Markuerkiaga et al., 2016).

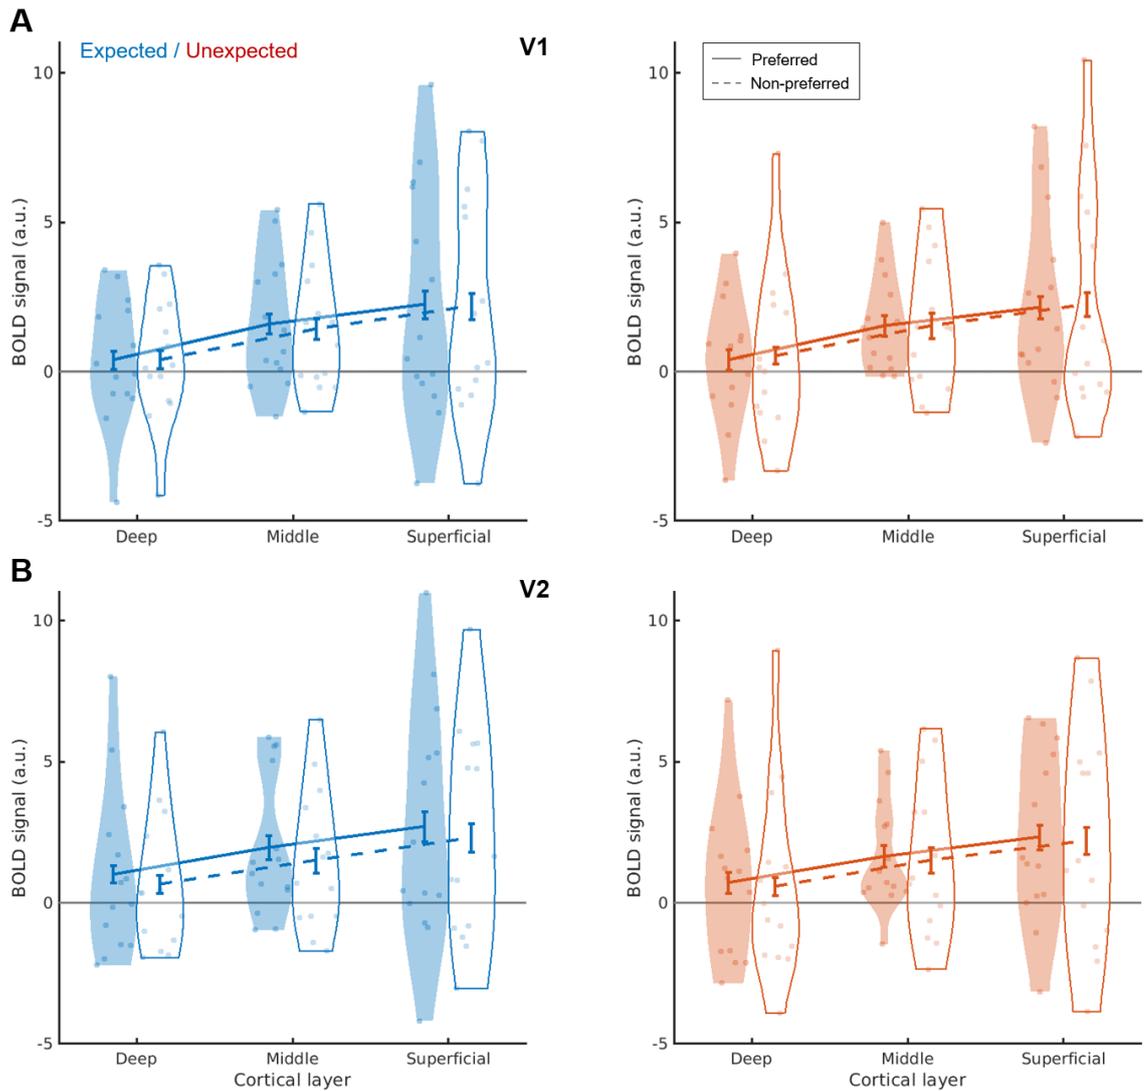


Figure 5.7. (A) BOLD responses in V1 across cortical layers for expected (left panel) and unexpected (right panel) events. Within each layer, responses are separated according to voxels preferring (solid lines) or non-preferring (dashed lines) the presented stimulus orientation, and demonstrates global BOLD signal before orientation-specific subtractions. Violin plots show individual data points for preferring (filled in) and non-preferring (outlined) voxels across layers (deep, middle and superficial). Error bars indicate within-subject SEM.

In sum, these findings reveal trends that orientation-specific BOLD responses in V1 and V2 were larger on average for expected events than unexpected events across all layers. This trend is possibly consistent on a numerical level with sharpening accounts and the

previous neural effects observed in this thesis – because larger orientation-specific values reflect activity increases in voxels tuned towards the presented stimulus relative to voxels tuned away from it. This would suggest that there was a greater representation of the presented orientation when it was expected relative to unexpected. Interestingly, these patterns remained constant across layers which preliminarily suggests that there may not be distinct P and PE populations residing in different cortical layers, inconsistent with the central principle of predictive coding theories (Bastos et al., 2012). However, such conclusions would be highly premature at this stage due to the lack of power and several methodological and analytical confounds that could account for these patterns – some of which will be described in the following section.

Examination of the univariate BOLD signal across layers demonstrates patterns consistent with known properties of the laminar-specific BOLD signal, which increases from deep to superficial layers as a function of venous blood draining to the pial surface (Duvernoy et al., 1981; Koopmans et al., 2010; Markuerkiaga et al., 2016). Moreover, equivalent orientation-specific activation across layers likely reflects propagation of signals due to laminar communications when both top-down and bottom-up processes are present, similarly to has been found in recent literature for expected events (Aitken et al., 2020). Introducing analysis techniques that compare orientation-specific signals – where draining is thought to have a more general influence across the column and therefore be equivalent for CW and CCW preferring voxels – and spatial regression approaches to de-correlate signals across layers are thought to reduce limitations such as these in laminar fMRI. Reducing such influences was particularly important in this experiment because it presented unexpected events for the first time, and therefore presenting these events as bottom-up input was necessary for examining such processing. Despite these mitigations, it is possible that such techniques may not have

been successful enough at disentangling signals across layers when the effects of interest in this Experiment are opposing in deep and superficial layers. Specifically, if signals are carried up the cortical column from deep to superficial layers, the expectation enhancement effects observed across all layers might instead reflect the hypothesised deep layer effects that are propagated upwards and ‘swamp’ any alternative processes in superficial layers evoked by unexpected events. Due to this possibility, it would be difficult to conclude from these results that there are not distinct processes operating in different cortical layers when expected and unexpected events are presented.

Furthermore, it is noticeable that orientation-specific numerical patterns in this experiment tend to peak and demonstrate the largest expectation effects in the middle layers across ROIs. This numerical effect is curious because feedback projections that purportedly convey top-down processes are known to have dense terminations in deep and superficial layers while avoiding middle layers (Anderson & Martin, 2009; Rockland & Virga, 1989; van Kerkoerle et al., 2017), and may be more consistent with findings that bottom-up input targets the middle layers of lower-level regions (Hubel & Wiesel, 1972; Lawrence, Norris, et al., 2019). However, this experiment hypothesised that there would be no orientation-specific differences between expected and unexpected conditions in middle layers because of this known functional organisation of feedback connections. Perhaps even the opposite trend would be hypothesised – that unexpected events evoke larger orientation-specific responses – if processing of PE operates as assumed by predictive coding theories, because feedforward connections terminating in the middle layers are thought to convey the mismatch, or error, between the input and predictions (Bastos et al., 2012) and unexpected events would indeed elicit greater error. Such processes may not have been observed in the present experiment

however if the effects in middle layers reflect contributions from signals carried up from expectation enhancement effects in deep layers.

Furthermore, some analytical problems at the boundary-based registration phase of the analysis pipeline are present in this data, whereby the boundary between GM and CSF is often incorrectly estimated, and therefore it is possible that the co-registration of these boundaries to functional scans does not accurately capture the most active voxels for the presented stimuli. Such issues are likely due to the novelty of 7T fMRI methods and analysis techniques in the field that are continuously being developed. Since this experiment was one of the first experiments to run at WCHN (UCL) on this scanner, there is ongoing work to optimise the sequences used in acquiring the EPI data here, which often contained some large posterior artefacts that overlapped with the ROIs. These issues, along with a lack of power, likely contribute to the large variability in effects across layers and therefore to the null effects overall. Due to these possibilities, it would be premature to assume that alternative processes do not operate across layers, and instead only very preliminary conclusions can be made about the enhancing nature of action predictions in deep layers of visual regions, considering they are the layers least confounded by these issues raised.

5.3. General Discussion

This chapter aimed to examine the possibility that perceptual predictive processes are determined by distinct populations of neurons representing P and PE in different cortical layers, and therefore whether such a distinction could help to reconcile findings reporting evidence for both upweighting and downweighting predictive processes.

Experiment 7 first examined the extent to which degradation of contingencies between action and outcomes influenced perception. It found that RT and accuracy benefits for

expected events were directly modulated by the degradation of contingency at test. The highest contingency level (75%) demonstrated the largest RT expectation benefits and was therefore adopted in Experiment 8 in the hope of maximising the predictive neural effects for examining the laminar profile of unexpected events for the first time in V1 and V2. Although no significant effects were observed across layers, likely due to power and analytical issues, numerical trends demonstrated that there was a larger orientation-specific response for expected, than unexpected, events across all layers and both ROIs. This pattern, especially in the deep layers, is consistent with previous findings in the literature demonstrating that predictions, or hypotheses, about our upcoming sensory input are encoded in the deep layers of early visual regions (Aitken et al., 2020; Kok, Bains, Van Mourik, et al., 2016; Lee & Nguyen, 2001). The effects might also be consistent with sharpening mechanisms that have been observed to underlie non-laminar effects of prediction in the same regions (Experiment 6; Kok, Jehee, et al., 2012; Yon et al., 2018).

However, as discussed above (see section 5.2.2.), levels of orientation specific signal that did not significantly differ across layers could reflect propagation of signal when top-down and bottom-up processes interact. This is particularly relevant for Experiment 8 because it manipulated prediction of sensory events always in the presence of bottom-up input. It is unclear whether such findings are the result of imperfect analysis techniques that are vulnerable to venous blood drain influences carrying signal from deep layers to middle and superficial layers, or whether it could simply reflect communication across layers when bottom-up input is present and the haemodynamic response function of the BOLD signal examines a relatively long time window. In support of the former possible explanation, studies have demonstrated that Gradient-Echo (GE) BOLD is much more susceptible to blood draining influences than other

BOLD measures such as Spin-Echo (Markuerkiaga et al., 2016) and 3D Gradient and Spin-Echo (GRASE; Moerel et al., 2018; Muckli et al., 2015). However, studies using comparable analysis techniques to in the present study (with GE BOLD, looking at orientation-specific signals) have demonstrated laminar specific influences of top-down and bottom-up process that differ across layers, and are thereby thought to be able to minimise such influences (Aitken et al., 2020; Lawrence et al., 2018; Lawrence, Formisano, et al., 2019). The activation across layers for expected events is comparable to findings that also presented bottom-up input and observe orientation-specific activation that does not differ across layers (Aitken et al., 2020). However, an explanation that propagation across layers reflects communications between layers when bottom-up and top-down processes interact could be likely under predictive coding assumptions, given that P and PE populations are thought to interact to minimise or ‘explain away’ errors of prediction violation and update future hypotheses, at each level of the hierarchy (Bastos et al., 2012; Friston, 2005).

A further, but non-laminar-related, possibility that could explain the high variability across participants’ data in Experiment 8 and thus the null effects, may relate to the differences in training experienced preceding the scanning session. In particular, participants in all other experiments in this thesis (Chapters 2-4) completed a training session under controlled experimental conditions where both training and test sessions were contextually identical. Due to COVID-19 restrictions, training sessions in Experiments 7 & 8 were completed online and therefore the exact environmental conditions in which these sessions were completed cannot be precisely determined. This might be problematic because insufficient concentration, or high environmental distractibility, during the task may have hindered learning about the association between action and stimulus orientation and therefore reduced predictive effects at test. On the

other hand, Experiment 7 suggested that reliable perceptual effects could be observed on the basis of such online training sessions, however one potentially key difference to the protocol of Experiment 8 was that training and test sessions both occurred online and therefore the contextual environment likely remained stable across sessions. Experiment 8 instead presented contextually different training and test sessions, given that training was completed online and therefore in participants' homes, while test sessions occurred in the scanner. This difference might have been problematic if perceptual effects of learning are dependent upon the context in which they have been learned (Bouton & King, 1983; Cook et al., 2012).

This possibility could also explain the lack of significant differences for RT or accuracy between expected and unexpected trials, yet Experiment 7 demonstrates large RT and accuracy benefits for expected events with the same contingency at test. A difference that is perhaps more plausible however, is that Experiment 7 was more optimally designed for demonstrating influences of expectation via reaction times in a way that Experiments 6 and 8 – that show no behavioural benefits – were not. In particular, participants were able to respond as soon as the stimulus appeared in Experiment 7, while ISIs delayed participants' responses in Experiments 6 and 8 to ensure response activity could not correlate with stimulus activity. Additionally, in Experiment 7 the response question was blocked – whereas in Experiments 6 and 8 it was randomised on each trial – such that participants could more easily anticipate their response on the basis of the action to-be-performed in the former. These differences therefore make comparisons across behavioural results difficult at this stage.

Nevertheless, this chapter highlights the importance for future work to optimise methodological and analytical approaches for laminar fMRI experiments in order to

provide reliable evidence for separable influences of P and PE in different cortical layers when presenting expected and unexpected events. Such investigations will also help to determine whether predictive coding accounts of sensory processing across layers can help to resolve the range of neural paradoxical findings of prediction within sensory cortical regions.

Chapter 6: General Discussion

6.1. Thesis summary

This thesis sought to investigate how action predictions influence perception. In action domains, perception of expected action outcomes is thought to be attenuated, or ‘cancelled’ (Blakemore et al., 1998; Dogge, Custers, & Aarts, 2019; Fiehler et al., 2019; Kilteni & Ehrsson, 2017a). Such a mechanism is proposed to be adaptive because surprising inputs are more useful, e.g., signalling the need to take new courses of action and update models of the world. However, theories outside of action and in broader sensory cognition domains purport that predicted events are instead perceptually facilitated – a process that is thought to be adaptive in order to enable us to generate largely veridical representations of our noisy sensory world (de Lange et al., 2018; Kersten et al., 2004). There has been an array of literature purportedly supporting both theories. However, studies tend to adopt distinct manipulations across domains and therefore comparison is difficult, such that it would be premature to conclude that action predictions indeed influence perception differently. The experiments in this thesis aimed to determine whether action predictions exert a distinct downweighting influence on perception and sensory processing or whether the influences are comparable to those seen outside of action when adopting similar manipulations.

Chapters 2 and 3 examined whether sensorimotor prediction necessarily attenuates perception of tactile events, as is widely assumed in the action literature, or whether predicted events may be perceptually upweighted – in line with theories in the wider sensory cognition literature. Chapter 2 adopted a force judgement paradigm frequently used in the action literature to demonstrate tactile attenuation, and included an additional condition to investigate the role of active contact during tactile force

judgements to determining effects (Experiment 1). Experiment 1 replicated typical findings that self-produced forces are rated as less intense than similar forces presented in the absence of action. However, this effect was reversed when the active finger did not make contact with any object, instead demonstrating enhanced force judgements during action. In principle, the attenuation effect in the contact condition may be generated by biasing due to gating of the stimulation on the active effector, but strong conclusions cannot be drawn here given that an active-passive comparison confounds examination of a number of potential mechanisms. More importantly, the finding of enhanced force judgements when this contact was removed could be consistent with upweighting theories suggesting that perception of expected events is enhanced. However, such conclusions could not be made before manipulating statistical probabilities between action and outcomes to determine predictive influences of action on tactile perception.

To this end, and to enable empirical comparisons with the broader sensory cognition literature, Chapter 3 presented three experiments that adopted a version of the force judgement paradigm but combined it with a manipulation of prediction according to learned statistical regularities. In these experiments, participants were trained on the association between one of two action types and the location of one of two tactile forces. Experiments 2 and 3 revealed that expected tactile action outcomes are perceived more, not less, intensely than those that were unexpected on the basis of statistical learning. Additional computational modelling in Experiment 3 suggested that expectations altered the way sensory evidence was accumulated – increasing the gain afforded to expected tactile signals. Finally, Experiment 4 examined a hypothesis that both up- and downweighting process operate during perceptual experience but at different time points (Press et al., 2020b), by presenting tactile vibrations of different

durations. Experiment 4 revealed similar upweighting influences of expectation on tactile perception for short duration events (in line with Experiments 2 and 3), but such influences were abolished when longer duration tactile forces were presented – supporting the possibility of opposing influences at different time points. However, strong conclusions were not drawn here about a null effect.

Taken together, these findings are consistent with upweighting perceptual accounts from outside of action domains, proposing that prior expectations are combined with sensory evidence to generate more veridical perceptual interpretations of our noisy environment – thereby rendering expected events more intense. The present findings indicate that these upweighting mechanisms operate similarly in action domains when predictions are manipulated in the same way, and are harder to reconcile with prominent downweighting theories from action (Blakemore et al., 1998; Dogge, Custers, & Aarts, 2019; Fiehler et al., 2019; Kiltner & Ehrsson, 2017a), which propose that expected action outcomes are perceptually attenuated.

Chapter 4 examined whether action predictions influence early visual processing in a distinct fashion, again examined via learned associations between actions and visual outcomes. Experiment 5 presented a perceptual experiment to investigate the nature of sensorimotor prediction on visual orientation discrimination, as a pilot for Experiment 6 that adopted fMRI techniques. Participants were trained to learn the relationship between action and the orientation of grating stimuli (clockwise or counter-clockwise oriented), and performed an orientation discrimination task at test. Although this experiment revealed no perceptual biasing in PSE measures, it revealed an RT benefit for expected events. The absence of the perceptual biasing effect suggests that action expectations may not generate ‘peak shift’ effects in a way that has been found once

before in the wider sensory cognition literature. Nevertheless, the presence of the RT benefit suggested that participants were able to learn to expect visual orientations on the basis of action, which allowed this task to be used for a subsequent fMRI study (Experiment 6). Experiment 6 investigated how action predictions influence V1 activity associated with expected and unexpected stimuli. Again, the primary aim was to determine whether sensorimotor prediction influences sensory processing in a distinct fashion from prediction mechanisms outside of action. Experiment 6 demonstrated using MVPA that expected action-outcomes could be decoded with superior accuracy than unexpected events, and univariate analyses demonstrated that expectation suppression was only observed in voxels tuned away from the presented orientation. These results support sharpening theories of perceptual processing and are consistent with studies outside of action that reveal similar sharpening mechanisms underlying prediction, and therefore suggest that sensorimotor prediction mechanisms do not operate in a distinct neural fashion from other types of prediction.

Finally, Chapter 5 aimed to determine whether discrepant findings within the neural literature could be reconciled by a conclusion that distinct neural populations representing predictions and prediction errors reside in different cortical layers.

Experiment 7 first investigated whether the contingency between learned action-outcome associations modulates behavioural RT benefits for expected events, and that is indeed what was found. The highest contingency group demonstrated the largest expectation benefit, and these effects reduced with degrading contingencies between action and visual outcomes. Experiment 8 adopted this paradigm using the highest contingency level to maximise potential neural predictive effects and to examine the laminar activity profile of expected and unexpected events in visual cortical regions. No concrete conclusions could be drawn from this experiment due to a lack of power and

possible analytical flaws, however trends in the data indicated that orientation-specific activity for expected events was greater across all layers than for unexpected events.

6.2. How could the perceptual prediction paradox be reconciled?

This thesis has demonstrated influences of action prediction on perception and sensory processing in early visual regions that are comparable to influences of prediction outside of action and consistent with upweighting theories. Such accounts propose that it is adaptive for us to prioritise perception of expected events, generating largely accurate representations of the sensory environment and aiding veridicality. The findings are however inconsistent with cancellation theories in action, that propose that prioritising perception of unexpected events enables us to remain maximally sensitive to informative sensory events so that we can rapidly update our models of the world and perform corrective actions. It is therefore essential to consider how findings in support of these theories could be reconciled with the present findings. There are a variety of alternative processes that may be able to account for the discrepancy across literatures in action and sensory cognition domains that require assuming that non-predictive processes contribute to – or even generate – apparent cancellation effects during action. However, it is also possible that both predictive upweighting and downweighting processes operate in perception but under different contextual conditions – where differences in methodological approaches might mean that studies only capture one ‘winning’ influence and are not optimally designed to measure both processes. This section will expand on both of these possibilities.

6.2.1. Sensory gating

Experiment 1 replicated typical force judgement findings that self-produced forces are perceived less intensely than the same externally-generated forces presented in the

absence of action. Interestingly, this effect was reversed when no active contact occurred during action, but strong conclusions cannot be made regarding the cause of this flip due to a number of possible explanations of active-passive differences. One such explanation could relate to the fact that studies demonstrating tactile attenuation during action assume that the sensory events which coincide with action are more predicted when adopting active vs passive comparisons, and support cancellation theories that predictions about the sensory consequences of action, or ‘efference copy’, is subtracted from the resulting percept (Wolpert & Flanagan, 2001). However, it is not only predictive processes that could differ when comparing perception of events presented during action and when passive – as is the predominant manipulation performed in the action domain (e.g., Bays et al., 2005, 2006; Blakemore et al., 1998; Kiltner et al., 2019; Reznik et al., 2015; Shergill et al., 2003; Stanley & Miall, 2007; Wolpe et al., 2018). It is widely acknowledged that performing actions generates identity-general gating of sensation on a moving effector, but which can also be present on passive effectors in close proximity to the moving finger (Williams et al., 1998; Williams & Chapman, 2000). These effects are not thought to represent predictive processes because they have been observed at the earliest relay in the spinal cord (Seki & Fetz, 2012). It is plausible that such influences contribute to attenuation findings in active relative to passive conditions. Attenuation effects that reflect reduced perception during active vs passive conditions can also be observed in a variety of other sensory species, but are often attributed to prediction processes without demonstrating specific influences of prediction (e.g., Crapse & Sommer, 2008; Enikolopov et al., 2018; Schneider et al., 2018).

In humans, it has been acknowledged that action studies examining tactile perception are vulnerable to conflating both identity-general gating and identity-specific prediction

processes as a single attenuation mechanism (e.g., Gertz et al., 2017; Juravle & Spence, 2011; see Kilteni & Ehrsson, 2020b). Indeed, studies have aimed to dissociate the two influences in typical tactile force judgement paradigms, with manipulations of active contact (Bays et al., 2006) as well as movement in the typically-passive left hand receiving stimulation (Kilteni & Ehrsson, 2020b). As discussed in section 2.2., the former does not provide convincing evidence for such a dissociation since active vs passive manipulations are still adopted and it is unclear why active contact is a necessary requirement for actions to be predictive of sensory consequences. A very recent study by Kilteni and Ehrsson (2020b) however adopted an alternate approach and manipulated movement on both the finger delivering and receiving the tactile force. Typical ‘attenuation’ conditions therefore reproduced the typical force-judgement set-up whereby an active right finger reaches towards a passive left finger and a tactile force is mechanically applied to the left finger. However, forces could also be presented to a moving left finger (moving via a sliding platform) while participants’ right-hand remained passive (‘gating’ condition) or in an active condition where participants made an action towards the moving left finger to produce the force (‘attenuation + gating’ condition). Notably, these conditions were further compared against a baseline condition in which no movement occurred on either finger delivering or receiving tactile forces. The authors report typical tactile attenuation effects in the classic conditions (attenuation vs baseline), yet judgements in both conditions that included a ‘gating’ component (that involved forces presented to a moving left hand) were further reduced in both active and passive trial types relative to the equivalent conditions when the left hand did not move. The authors conclude that the additional reduction in forces across conditions that involved the ‘gating’ component on the left hand provides evidence for a separable mechanism, since typical attenuation effects were observed orthogonally (e.g.,

in active vs. passive conditions) to reduction in perceived intensity produced by gating (e.g., by left hand sliding platform).

However, it is difficult to draw conclusions about separable influences of predictive attenuation from attenuation produced by gating on the basis of these results. Moreover, gating mechanisms were manipulated on the left force-receiving hand and so it is difficult to determine how additive effects of suppression relating to movement of the two effectors implies that one of the effects is predictive. One possible explanation of the reduced force judgements in the 'attenuation + gating' condition, relative to the original 'attenuation' condition, could simply reflect a joint suppressive contribution of gating on both right and left hands during movement, whereby such contributions are only provided by the right hand in the typical 'attenuation' condition. It therefore still remains unclear as to whether 'predictive' tactile attenuation in these types of procedure is mediated by prediction mechanisms. Emphasis therefore needs to be placed on adopting paradigms in future studies that compare perception of events that differ in their levels of expectedness but always in the presence of action, to control for non-predictive influences that might influence perception during action.

Experiments 2-4 aimed to remove such influences in this way, firstly by recording actions via a motion tracker that did not require active contact, but more importantly by examining predictive influences according to statistical learning that enable comparisons between expected and unexpected events always in the presence of action. It is important to note that such manipulations reduce the possibility of gating generating observed differences but also control for a wide variety of other processes that likely differ during active and passive perception (Press & Cook, 2015). For example, in active perception conditions participants have a dual-task relative to the

single task in passive perception conditions, and classic accounts of working memory would predict fewer resources for processing the stimulus in the presence of such additional tasks (Baddeley, 1996; Press et al., 2020a).

6.2.2. Repetition suppression

As highlighted, there is a distinct lack of evidence that tactile attenuation is underpinned by predictive processes defined according to statistical learning, but rather via the assumption that tactile events that are self-produced should be more predictable than events that are externally produced. While non-predictive mechanisms such as sensory gating could induce attenuating influences on perception, it is possible that other attenuation effects could be generated by mechanisms shaping perception according to event repetition – given that repetition is frequently confounded with expectation (Feuerriegel et al., 2020). Repetition suppression is a term given to a phenomenon that demonstrates attenuated perception of repeated stimuli (Ofen et al., 2007) as well as reduced sensory processing associated with those stimulus features (Grill-Spector et al., 2006). Whether there is a distinction between mechanisms underlying repetition suppression and expectation suppression is a topic of much debate. For instance, it has been suggested that repetition suppression partly reflects influences of prediction because stimuli can be expected when they are continuously repeated (Grill-Spector et al., 2006; Summerfield et al., 2008; Summerfield & de Lange, 2014), while others have suggested that it may more likely reflect low-level adaptation (Carandini, 2000; Webster, 2015), or a combination of the two (Larsson & Smith, 2012).

A popular study cited in support of an expectation account, used fMRI to compare BOLD signal differences between face stimuli that could be either repeated or alternated (so that the presented face was followed by a novel face; Summerfield et al., 2008). The

authors reported a larger reduction in neural response when stimulus repetitions were likely, and therefore more expected, than when repetitions were unlikely and relatively unexpected – suggested to provide evidence for interacting processes. However, a study that used MEG to measure neural responses when repetition and prediction were orthogonalised by sometimes rendering stimulus alternations predictable, revealed distinct neural time-courses exhibited for repetition and prediction effects (Todorovic & De Lange, 2012). A further study adopting a similar paradigm but measured neuronal activity in macaques and found that stimulus repetition probability does not differentially modulate the suppression observed, and suggest that repetition suppression is not likely to rely on perceptual expectation influences (Kaliukhovich & Vogels, 2011). More recently, evidence from computational modelling (Alink et al., 2018) has emerged that repetition neural effects may be mediated by mechanisms functionally distinct from prediction effects – but with continued debates surrounding whether the dissociation might come from prediction effects that emerge via neural sharpening (Kok, Jehee, et al., 2012; Yon et al., 2018) but with repetition effects emerging via local neural scaling (Alink et al., 2018; but see Alink et al., 2020; Ramírez & Merriam, 2020).

It will therefore be important to determine to what extent repetition suppression relies on predictive processes and whether neural suppression effects reported in the literature might conflate both prediction and repetition processes. It is also essential to consider how repetition effects might influence behavioural results of sensory attenuation in action domains, because it is notable that recent paradigms that have claimed to be isolating prediction in fact frequently manipulate repetition alongside prediction. For example, a study measuring tactile intensity judgements on an active hand required participants to make actions towards an object and pick it up, while the mass

distribution of the object could be predictable or unpredictable in separate blocks (Voudouris et al., 2019). Importantly, predictable mass distributions were always the same in predictable blocks, yet in unpredictable blocks the mass distribution changed every trial. The study therefore reports stronger tactile attenuation in predictable blocks, however it is uncertain what mechanism could be responsible for these effects considering that equivalent mass distributions were experienced continuously in predictable blocks – increasing the likelihood of repetition effects – yet these distributions of mass were alternated every trial in unpredictable blocks.

Furthermore, as discussed in section 3.5., a tactile action study conducted by Kilteni et al. (2019) that aimed to determine whether tactile attenuation reflected predictive mechanisms according to statistical probabilities could be sensitive to similar influences. The authors trained participants that sensory events would be delayed by 100 ms, but only this one repeated mapping was experienced before they were then tested on events at that delay or at no delay (in different conditions) – finding a typical attenuation effect for delayed (but expected) events. However, in conditions where the tested delay and repeated delay matched, it is possible that repeated presentation of the 100 ms delay could explain attenuated perception rather than prediction of it, considering that every response trial was preceded by an extra five repeated training trials. It is therefore important for future studies examining predictive processes according to statistical regularities to take influences of repetition suppression into account, ensuring that there are multiple action-outcome associations presented throughout testing, reducing influences of continuous repetitions of one condition or trial type.

The experiments presented in this thesis aimed to adopt such procedures to reduce potentially competing repetition influences on perception. Namely, participants were trained on the relationship between one of two actions and one of two resulting sensory outcomes and the specific association presented on each trial was randomised. This meant that prediction influences were not examined in contexts where instances of action and sensory events were repeated or blocked, and this likely reduced influences of repetition suppression while enabling prediction mechanisms to be measured through statistical learning.

6.2.3. Expectation and signal strength

A further notable difference across literatures purporting support for both theories is that the evidence supporting cancellation tends to come from action paradigms examining suprathreshold (tactile) perception (e.g., Bays et al., 2005; Blakemore et al., 1998; Kiltner & Ehrsson, 2017b), whereas evidence supporting upweighting in sensory cognition domains often examines (visual) perceptual sensitivity (e.g., Chalk et al., 2010; Cheadle et al., 2015; Wyart et al., 2012). This distinction makes comparisons across studies and domains difficult, because it is possible that differences in stimulus strength determines empirical differences. Thus, although some of the downweighting effects reported in the action literature may have alternative explanations as discussed, it is also conceivable that a difference could often relate to the type of stimulus presented. Indeed, the opposing process theory (Press et al., 2020b) suggests that two processes operate to shape perception and these processes exhibit different influences that are dependent upon the extent of unexpectedness of a signal. Specifically, the theory posits that when there is low deviation between what we expect and the input, the secondary process highlighting the unexpected will not be produced. This will more likely be the case with weak sensory evidence, and thus with unexpected signals in signal detection

paradigms. Such events that generate low error likely do not demonstrate a need for model updating and when unexpectedness could be attributable to sensory noise.

However, when we perceive an unexpected event that is strong and therefore generating high surprise, later reactive mechanisms will highlight perception of this informative input to aid model updating (Press et al., 2020b). Therefore, the important assumption from this theory is that mechanisms highlighting unexpected events will only operate if the event generates high surprise, and therefore such processes are unlikely to be observed in high noise environments because unexpected events are not surprising enough.

Such processes could account for differences across domains that primarily adopt highly noisy detection paradigms and demonstrate evidence in favour of upweighting mechanisms, while studies that primarily present supra-threshold sensory events observe downweighting. Some recent studies have indeed provided evidence but within experimental set-ups for such a hypothesis. For example, perception of self-generated auditory events has been found to be attenuated relative to passive events when presented at suprathreshold intensity, while self-generated events presented at-threshold level were perceived louder than equivalent passive events (Paraskevoudi & SanMiguel, 2020; Reznik et al., 2015). It is therefore possible that the array of sensory attenuation effects on supra-threshold stimuli in the action literature could be accounted for by such a hypothesis, considering that surprise-based downweighting processes can only operate once highly surprising unexpected events can be easily detected.

This hypothesis is potentially at odds with the presented experiments of this thesis because suprathreshold events were always presented, but the results primarily demonstrate support for upweighting accounts. However, the observed tactile biases in

Experiments 1-4 might not be so inconsistent with claims that upweighting processes operate at the earliest perceptual time windows, considering that tactile events were presented for very short durations and therefore later processes would not have had the chance to influence perception before stimulus offset. Moreover, when long duration tactile events were presented in Experiment 4, upweighted effects were eradicated, suggesting the possibility that later surprise based mechanisms could have influenced perception – a process that could plausibly have occurred because events were highly detectable. Such conclusions concerning the influence of opposing processes in this experiment are however made tentatively because processing in early vs late time windows was not directly compared.

Sharpening effects were of course also observed in Experiment 6 (and tentatively in Experiment 8) when suprathreshold visual events tended to be of high contrast and presented for longer stimulus durations. Evidently, the fMRI methods adopted in these experiments cannot provide any information about the temporal specificity of these effects, so an interesting avenue for future research could adopt highly temporally specific methods such as EEG or MEG to investigate assumptions that opposing mechanisms might operate at different time points. Such future experiments could directly compare high and low noise settings to examine the possibility that downweighting effects predominate when highly surprising unexpected outcomes are presented, and could help to resolve some of the discrepancies across perceptual findings in the literature. This work should importantly be undertaken alongside laminar fMRI studies continuing the investigation into whether different neural effects predominate in different cortical laminae. All of these investigations must proceed alongside examinations of how neural processing specifically shapes perception –

bearing in mind that the theoretical paradox largely arises at the level of perception and the neural examinations are employed to inform this perceptual puzzle.

6.3. Expectation vs Attention

A further point to disentangle when determining predictive influences on perception will be the relationship between expectation and attention. As already discussed, while expectation influences perception according to statistical likelihood of event occurrence, attention prioritises perceptual information according to task relevance (Summerfield & Egner, 2009, 2016). However, in our natural environment as well as the majority of experimental paradigms, more probable events are also more relevant for task performance (e.g., Posner et al., 1980; Rohenkohl et al., 2014). As well as determining separable influences of likelihood and task relevance of presented stimuli, another factor that may be crucial for determining underlying mechanisms concerns task relevance of the learned associations.

Specifically, Experiment 6 and others (Aitken et al., 2020; González-García & He, 2021; Kok, Jehee, et al., 2012; Yon et al., 2018) have demonstrated stimulus-specific sharpening of expected representations, while others have revealed that predictable objects evoke stimulus-specific suppression in the same areas (Han et al., 2019; Richter et al., 2018; Richter & de Lange, 2019). A subtle difference between these studies reporting effects of expectation, even when they do not manipulate attention explicitly, could relate to whether the predictive associations themselves are task relevant. For example, the study by Richter et al. (2018) manipulated statistical learning of object pairs and used fMRI to measure neural response to expected and unexpected objects – yet participants were only required to respond when objects were presented upside-down, rendering these associations task-irrelevant. Therefore, expectations were task-

irrelevant in Richter et al.'s (2018) experiment in comparison to task-relevant predictions in the current study and those of Yon et al. (2018) and Kok et al. (2012) demonstrating sharpening. There are indeed ongoing debates about exactly how expectation and attention interact in sensory processing, where some studies have indicated that attention can reverse suppression effects (Kok, Rahnev, et al., 2012), yet others have demonstrated suppression effects only when predicted features are task relevant (Richter & de Lange, 2019).

The experiments presented in this thesis primarily manipulated statistical probabilities in attempt to orthogonalise the influences of expectation and attention but did not manipulate task relevance explicitly. Interestingly, similar conclusions regarding the operation of upweighting mechanisms could be drawn across the experiments in this thesis regardless of whether trained associations were task-relevant (Experiments 5-8) or task-irrelevant (Experiments 1-4). Nevertheless, it will be essential for future work to orthogonalise expectation and attention in determining the relative level of overlap or interaction in generating upweighting or downweighting influences on perception.

6.4. Implications for wider applications of Cancellation theories in action

Regardless of the ultimate resolution of this debate, the important conclusion from the present studies is that sensorimotor prediction does not appear to exhibit a qualitatively distinct influence on perception – demonstrating the need to update theoretical accounts to incorporate the vast array of data revealing both up- and downweighting across perceptual disciplines. This will be crucial for determining the typical influence of prediction on perception in healthy young populations, as well as older and clinical populations. Tactile attenuation during action – and even more specifically, the particular force judgement paradigm – has been used to demonstrate sensory differences

associated with healthy ageing (Wolpe et al., 2016), motor severity in Parkinson's disease (Wolpe et al., 2018) and hallucinatory severity in schizophrenia (Shergill et al., 2014). For example, studies have demonstrated that patients with schizophrenia tend to show reduced perceptual tactile attenuation effects (Shergill et al., 2005) and exhibit less attenuated responses in secondary somatosensory cortices which can be predicted by hallucinatory severity (Shergill et al., 2014). Furthermore, the assumption that action predictions have a particular and distinctive attenuating influence on perception has been used to advance a number of theoretical accounts, including those of hallucinations in clinical disorders such as schizophrenia and diseases such as Parkinson's and Alzheimer's (Corlett et al., 2019). Such theories typically attribute differences in sensory processing to aberrant prediction of action consequences, where hallucinations arise due to a failure to attenuate those expected consequences. However, if predictions shape perception similarly regardless of domain then these theories may need revisiting.

6.5. Conclusions

To conclude, the findings of this thesis suggest that sensorimotor prediction can increase the perceived intensity of tactile events and sharpen sensory representations of expected visual action outcomes. These findings are in line with upweighting predictive accounts that likely enable us to generate largely veridical experiences in the face of sensory noise, and suggest that sensorimotor prediction influences perception via qualitatively similar mechanisms to other types of prediction outside of action. Future work must determine how we generate both veridical and informative experiences on the basis of our expectations, but this thesis suggests that solutions likely apply similarly regardless of domain.

References

- Aitken, F., Menelaou, G., Warrington, O., Koolschijn, R. S., Corbin, N., Callaghan, M. F., & Kok, P. (2020). Prior expectations evoke stimulus-specific activity in the deep layers of the primary visual cortex. *PLoS Biology*, *18*(12), e3001023. <https://doi.org/10.1371/journal.pbio.3001023>
- Albers, A. M., Meindertsma, T., Toni, I., & de Lange, F. P. (2018). Decoupling of BOLD amplitude and pattern classification of orientation-selective activity in human visual cortex. *NeuroImage*, *180*, 31–40. <https://doi.org/10.1016/J.NEUROIMAGE.2017.09.046>
- Alink, A., Abdulrahman, H., & Henson, R. N. (2018). Forward models demonstrate that repetition suppression is best modelled by local neural scaling. *Nature Communications*, *9*(1), 1–10. <https://doi.org/10.1038/s41467-018-05957-0>
- Alink, A., Abdulrahman, H., & Henson, R. N. (2020). Reply to ‘Forward models of repetition suppression depend critically on assumptions of noise and granularity.’ *Nature Communications*, *11*(1), 1–3. <https://doi.org/10.1038/s41467-020-18316-9>
- Alink, A., Schwiedrzik, C. M., Kohler, A., Singer, W., & Muckli, L. (2010). Stimulus predictability reduces responses in primary visual cortex. *Journal of Neuroscience*, *30*(8), 2960–2966. <https://doi.org/10.1523/JNEUROSCI.3730-10.2010>
- Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R., & Kievit, R. A. (2019). Raincloud plots: A multi-platform tool for robust data visualization [version 1; peer review: 2 approved]. *Wellcome Open Research*, *4*, 63. <https://doi.org/10.12688/wellcomeopenres.15191.1>
- Anderson, J. C., & Martin, K. A. C. (2009). The synaptic connections between cortical areas V1 and V2 in macaque monkey. *Journal of Neuroscience*, *29*(36), 11283–11293. <https://doi.org/10.1523/JNEUROSCI.5757-08.2009>
- Anzai, A., Peng, X., & Van Essen, D. C. (2007). Neurons in monkey visual area V2 encode combinations of orientations. *Nature Neuroscience*, *10*(10), 1313–1321. <https://doi.org/10.1038/nn1975>
- Baddeley, A. (1996). Exploring the Central Executive. *The Quarterly Journal of Experimental Psychology Section A*, *49*(1), 5–28. <https://doi.org/10.1080/713755608>
- Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, *5*(8), 617–629. <https://doi.org/10.1038/nrn1476>
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical Microcircuits for Predictive Coding. *Neuron*, *76*(4), 695–711. <https://doi.org/10.1016/j.neuron.2012.10.038>
- Bays, P. M., Flanagan, J. R., & Wolpert, D. M. (2006). Attenuation of self-generated tactile sensations is predictive, not postdictive. *PLoS Biology*, *4*(2), 281–284. <https://doi.org/10.1371/journal.pbio.0040028>
- Bays, P. M., Wolpert, D. M., & Flanagan, J. R. (2005). Perception of the consequences of self-action is temporally tuned and event driven. *Current Biology*, *15*(12), 1125–

1128. <https://doi.org/10.1016/j.cub.2005.05.023>

- Blakemore, S. J., Frith, C. D., & Wolpert, D. M. (1999). Spatio-Temporal Prediction Modulates the Perception of Self-Produced Stimuli. *Journal of Cognitive Neuroscience*, *11*(5), 551–559. <https://doi.org/10.1162/089892999563607>
- Blakemore, S. J., Frith, C. D., & Wolpert, D. M. (2001). The cerebellum is involved in predicting the sensory consequences of action. *Neuroreport*, *12*(9), 1879–1884. <https://doi.org/10.1097/00001756-200107030-00023>
- Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience*, *1*(7), 635–640. <https://doi.org/10.1038/2870>
- Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (2000). Why can't you tickle yourself? *NeuroReport*, *11*(11), R11–R16. <https://doi.org/10.1097/00001756-200008030-00002>
- Blom, T., Feuerriegel, D., Johnson, P., Bode, S., & Hogendoorn, H. (2020). Predictions drive neural representations of visual events ahead of incoming sensory information. *Proceedings of the National Academy of Sciences*, *117*(13), 7510–7515. <https://doi.org/10.17605/OSF.IO/V9YX8.y>
- Bouton, M. E., & King, D. A. (1983). Contextual control of the extinction of conditioned fear: Tests for the associative value of the context. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*(3), 248–265. <https://doi.org/10.1037/0097-7403.9.3.248>
- Brodmann, K. (2007). Brodmann's: Localisation in the Cerebral Cortex. In *Springer Science & Business Media*. <https://www.appliedneuroscience.com/PDFs/Brodmann.pdf>
- Brown, H., Adams, R. A., Pares, I., Edwards, M., & Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive Processing*, *14*(4), 411–427. <https://doi.org/10.1007/s10339-013-0571-3>
- Campbell, K. L., Healey, M. K., Lee, M. M. S., Zimmerman, S., & Hasher, L. (2012). Brief report: Age differences in visual statistical learning. *Psychology and Aging*, *27*(3), 650–656. <https://doi.org/10.1037/a0026780>
- Carandini, M. (2000). Visual cortex: Fatigue and adaptation. *Current Biology*, *10*(16), R605–R607. [https://doi.org/10.1016/S0960-9822\(00\)00637-0](https://doi.org/10.1016/S0960-9822(00)00637-0)
- Cardoso-Leite, P., Mamassian, P., Schütz-Bosbach, S., & Waszak, F. (2010). A New Look at Sensory Attenuation: Action-Effect Anticipation Affects Sensitivity, Not Response Bias. *Psychological Science*, *21*(12), 1740–1745. <https://doi.org/10.1177/0956797610389187>
- Cardoso-Leite, P., Mamassian, P., Schütz-Bosbach, S., & Waszak, F. (2010). A New Look at Sensory Attenuation. *Psychological Science*, *21*(12), 1740–1745. <https://doi.org/10.1177/0956797610389187>
- Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, *7*(3), 308–313. <https://doi.org/10.1038/nn1194>

- Chalk, M., Seitz, A. R., & Seriès, P. (2010). Rapidly learned stimulus expectations alter perception of motion. *Journal of Vision*, *10*(8), 1–18. <https://doi.org/10.1167/10.8.2>
- Cheadle, S., Egner, T., Wyart, V., Wu, C., & Summerfield, C. (2015). Feature expectation heightens visual sensitivity during fine orientation discrimination. *Journal of Vision*, *15*(14), 14. <https://doi.org/10.1167/15.14.14>
- Christensen, A., Ilg, W., & Giese, M. A. (2011). Spatiotemporal tuning of the facilitation of biological motion perception by concurrent motor execution. *The Journal of Neuroscience*, *31*(9), 3493–3499. <https://doi.org/10.1523/JNEUROSCI.4277-10.2011>
- Christoffels, I. K., Formisano, E., & Schiller, N. O. (2007). Neural correlates of verbal feedback processing: An fMRI study employing overt speech. *Human Brain Mapping*, *28*(9), 868–879. <https://doi.org/10.1002/hbm.20315>
- Conrad, V., Vitello, M. P., & Noppeney, U. (2012). Interactions Between Apparent Motion Rivalry in Vision and Touch. *Psychological Science*, *23*(8), 940–948. <https://doi.org/10.1177/0956797612438735>
- Cook, R., Bird, G., Catmur, C., Press, C., & Heyes, C. (2014). Mirror neurons: From origin to function. *Behavioral and Brain Sciences*, *37*(2), 177–192. <https://doi.org/10.1017/S0140525X13000903>
- Cook, R., Dickinson, A., & Heyes, C. (2012). Contextual modulation of mirror and counter-mirror sensorimotor associations. *Journal of Experimental Psychology: General*, *141*(4), 774. <https://doi.org/10.1037/a0027561>
- Cook, R., Press, C., Dickinson, A., & Heyes, C. (2010). Acquisition of Automatic Imitation is Sensitive to Sensorimotor Contingency. *Journal of Experimental Psychology: Human Perception and Performance*, *36*(4), 840.
- Corlett, P. R., Horga, G., Fletcher, P. C., Alderson-Day, B., Schmack, K., & Powers, A. R. (2019). Hallucinations and Strong Priors. *Trends in Cognitive Sciences*, *23*(2), 114–127. <https://doi.org/10.1016/j.tics.2018.12.001>
- Crapse, T. B., & Sommer, M. A. (2008). Corollary discharge across the animal kingdom. *Nature Reviews Neuroscience*, *9*(8), 587–600. <https://doi.org/10.1038/nrn2457>
- Crawford, L. E., Huttenlocher, J., & Engebretson, P. H. (2000). Category effects on estimates of stimuli: Perception or Reconstruction? *Psychological Science*, *11*(4), 280–284.
- de Lange, F. P., Heilbron, M., & Kok, P. (2018). How Do Expectations Shape Perception? *Trends in Cognitive Sciences*, *22*(9), 764–779. <https://doi.org/10.1016/j.tics.2018.06.002>
- 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*, *34*(4), 1400–1410. <https://doi.org/10.1523/JNEUROSCI.1094-12.2013>
- Desantis, A., Mamassian, P., Lisi, M., & Waszak, F. (2014). The prediction of visual

- stimuli influences auditory loudness discrimination. *Exp Brain Res*, 232, 3317–3324. <https://doi.org/10.1007/s00221-014-4001-2>
- Dickinson, A. (1980). *Contemporary Animal Learning Theory*. Cambridge University Press.
- Dickinson, A., & Charnock, D. J. (1985). Contingency effects with maintained instrumental reinforcement. *The Quarterly Journal of Experimental Psychology*, 37(4), 397–416. <https://doi.org/10.1080/14640748508401177>
- Dogge, M., Custers, R., & Aarts, H. (2019). Moving Forward: On the Limits of Motor-Based Forward Models. *Trends in Cognitive Sciences*, 23(9), 743–753. <https://doi.org/10.1016/J.TICS.2019.06.008>
- Dogge, M., Custers, R., Gayet, S., Hoijsink, H., & Aarts, H. (2019). Perception of action-outcomes is shaped by life-long and contextual expectations. *Scientific Reports*, 9(1), 5225. <https://doi.org/10.1038/s41598-019-41090-8>
- Duffy, S., Huttenlocher, J., Hedges, L. V., & Crawford, L. E. (2010). Category effects on stimulus estimation: Shifting and skewed frequency distributions. *Psychonomic Bulletin and Review*, 17(2), 224–230. <https://doi.org/10.3758/PBR.17.2.224>
- Durrant, S. J., Taylor, C., Cairney, S., & Lewis, P. A. (2011). Sleep-dependent consolidation of statistical learning. *Neuropsychologia*, 49(5), 1322–1331. <https://doi.org/10.1016/j.neuropsychologia.2011.02.015>
- Duvernoy, H. M., Delon, S., & Vannson, J. L. (1981). Cortical blood vessels of the human brain. *Brain Research Bulletin*, 7(5), 519–579. [https://doi.org/10.1016/0361-9230\(81\)90007-1](https://doi.org/10.1016/0361-9230(81)90007-1)
- Ekman, M., Kok, P., & De Lange, F. P. (2017). Time-compressed preplay of anticipated events in human primary visual cortex. *Nature Communications*, 8(1), 1–9. <https://doi.org/10.1038/ncomms15276>
- Elsner, B., & Hommel, B. (2004). Contiguity and contingency in action-effect learning. *Psychological Research*, 68(2–3), 138–154. <https://doi.org/10.1007/s00426-003-0151-8>
- Enikolopov, A. G., Abbott, L. F., & Sawtell, N. B. (2018). Internally Generated Predictions Enhance Neural and Behavioral Detection of Sensory Stimuli in an Electric Fish. *Neuron*, 99(1), 135–146. <https://doi.org/10.1016/j.neuron.2018.06.006>
- Erickson, C. A., & Desimone, R. (1999). Responses of Macaque Perirhinal Neurons during and after Visual Stimulus Association Learning. *Journal of Neuroscience*, 19(23), 10404–10416.
- Ester, E. F., Sprague, T. C., & Serences, J. T. (2020). Categorical biases in human occipitoparietal cortex. *Journal of Neuroscience*, 40(4), 917–931. <https://doi.org/10.1523/JNEUROSCI.2700-19.2019>
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cerebral Cortex*, 1, 1–47.
- Feuerriegel, D., Blom, T., & Hogendoorn, H. (2021). Predictive activation of sensory

- representations as a source of evidence in perceptual decision-making. *Cortex*, *136*, 140–146. <https://doi.org/10.1016/j.cortex.2020.12.008>
- Feuerriegel, D., Churches, O., Coussens, S., & Keage, H. A. D. (2019). Temporal expectations modulate face image repetition suppression of early stimulus evoked event-related potentials. *Neuropsychologia*, *122*, 76–87. <https://doi.org/10.1016/j.neuropsychologia.2018.11.010>
- Feuerriegel, D., Keage, H. A. D., Rossion, B., & Quek, G. L. (2018). Immediate stimulus repetition abolishes stimulus expectation and surprise effects in fast periodic visual oddball designs. *Biological Psychology*, *138*, 110–125. <https://doi.org/10.1016/j.biopsycho.2018.09.002>
- Feuerriegel, D., Vogels, R., & Kovács, G. (2020). Evaluating the Evidence for Expectation Suppression in the Visual System. *PsyArXiv*. <https://doi.org/10.31234/OSF.IO/Y7639>
- Fiehler, K., Brenner, E., & Spering, M. (2019). Prediction in goal-directed action. *Journal of Vision*, *19*(9), 1–21. <https://doi.org/10.1167/19.9.10>
- Firestone, C., & Scholl, B. (2016). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioural and Brain Sciences*, *39*, e229. <https://doi.org/https://doi.org/10.1017/S0140525X15000965>
- Flounders, M. W., Gonzá Lez-García, C., Hardstone, R., & He, B. J. (2019). Neural dynamics of visual ambiguity resolution by perceptual prior. *Elife*, *8*, e41861. <https://doi.org/10.7554/eLife.41861.001>
- Friston, K. (2005). A theory of cortical responses. *Phil. Trans. R. Soc. B*, *360*, 815–836. <https://doi.org/10.1098/rstb.2005.1622>
- Friston, K. (2008). Hierarchical models in the brain. *PLoS Computational Biology*, *4*(11). <https://doi.org/10.1371/journal.pcbi.1000211>
- Fritsche, M., & de Lange, F. P. (2019). Reference repulsion is not a perceptual illusion. *Cognition*, *184*, 107–118. <https://doi.org/10.1016/J.COGNITION.2018.12.010>
- Gagnepain, P., Henson, R. N., & Davis, M. H. (2012). Temporal predictive codes for spoken words in auditory cortex. *Current Biology*, *22*(7), 615–621. <https://doi.org/10.1016/j.cub.2012.02.015>
- Gallace, A., Zeeden, S., Röder, B., & Spence, C. (2010). Lost in the move? Secondary task performance impairs tactile change detection on the body. *Consciousness and Cognition*, *19*(1), 215–229. <https://doi.org/10.1016/j.concog.2009.07.003>
- Gandolfo, M., & Downing, P. E. (2019). Causal Evidence for Expression of Perceptual Expectations in Category-Selective Extrastriate Regions. *Current Biology*, *29*(15), 2496–2500.e3. <https://doi.org/10.1016/J.CUB.2019.06.024>
- Gertz, H., Voudouris, D., & Fiehler, K. (2017). Reach-relevant somatosensory signals modulate tactile suppression. *Journal of Neurophysiology*, *117*(6), 2262–2268. <https://doi.org/10.1152/jn.00052.2017>
- González-García, C., & He, B. J. (2021). A Gradient of Sharpening Effects by Perceptual Prior across the Human Cortical Hierarchy. *Journal of Neuroscience*,

41(1), 167–178. <https://doi.org/10.1523/JNEUROSCI.2023-20.2020>

- Gregory, R. L. (1997). Knowledge in perception and illusion. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 352(1358), 1121–1128. <https://doi.org/10.1098/rstb.1997.0095>
- Greve, D. N., & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, 48(1), 63–72. <https://doi.org/10.1016/j.neuroimage.2009.06.060>
- Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: neural models of stimulus-specific effects. *Trends in Cognitive Sciences*, 10(1), 14–23. <https://doi.org/10.1016/j.tics.2005.11.006>
- Guo, J., & Song, J.-H. (2019). Action Fluency Facilitates Perceptual Discrimination. *Psychological Science*, 30(10). <https://doi.org/10.1177/0956797619859361>
- Han, B., Mostert, P., & de Lange, F. P. (2019). Predictable tones elicit stimulus-specific suppression of evoked activity in auditory cortex. *NeuroImage*, 200, 242–249. <https://doi.org/10.1016/j.neuroimage.2019.06.033>
- Han, B., & van Rullen, R. (2016). Shape perception enhances perceived contrast: evidence for excitatory predictive feedback? *Scientific Reports*, 6, 22944. <https://doi.org/10.1038/srep22944>
- Haynes, J. D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience*, 7(7), 523–534. <https://doi.org/10.1038/nrn1931>
- Hebart, M. N., Görden, K., & Haynes, J.-D. (2015). The Decoding Toolbox (TDT): a versatile software package for multivariate analyses of functional imaging data. *Frontiers in Neuroinformatics*, 8, 88. <https://doi.org/10.3389/fninf.2014.00088>
- Hinds, O. P., Rajendran, N., Polimeni, J. R., Augustinack, J. C., Wiggins, G., Wald, L. L., Diana Rosas, H., Potthast, A., Schwartz, E. L., & Fischl, B. (2008). Accurate prediction of V1 location from cortical folds in a surface coordinate system. *NeuroImage*, 39(4), 1585–1599. <https://doi.org/10.1016/j.neuroimage.2007.10.033>
- Hubel, D. H., & Wiesel, T. N. (1972). Laminar and columnar distribution of geniculocortical fibers in the macaque monkey. *The Journal of Comparative Neurology*, 146(4), 421–450. <https://doi.org/10.1002/cne.901460402>
- Hudson, M., Bach, P., & Nicholson, T. (2017). You Said You Would! The Predictability of Other's Behavior From Their Intentions Determines Predictive Biases in Action Perception. *Journal of Experimental Psychology: Human Perception and Performance*, 44(2), 320. <https://doi.org/10.1037/xhp0000451>
- Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why Do Categories Affect Stimulus Judgment? *Journal of Experimental Psychology: General*, 129(2), 220–241. <https://doi.org/10.1037/0096-3445.129.2320>
- Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience*, 9(5), 690–696. <https://doi.org/10.1038/nn1691>

- Jazayeri, M., & Movshon, J. A. (2007). Integration of sensory evidence in motion discrimination. *Journal of Vision*, 7(12), 1–7. <https://doi.org/10.1167/7.12.7>
- Juravle, G., & Spence, C. (2011). Juggling reveals a decisional component to tactile suppression. *Experimental Brain Research*, 213(1), 87–97. <https://doi.org/10.1007/s00221-011-2780-2>
- Kaliukhovich, D. A., & Vogels, R. (2011). Stimulus repetition probability does not affect repetition suppression in macaque inferior temporal cortex. *Cerebral Cortex*, 21(7), 1547–1558. <https://doi.org/10.1093/cercor/bhq207>
- Kanai, R., Komura, Y., Shipp, S., & Friston, K. (2015). Cerebral hierarchies: Predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1668). <https://doi.org/10.1098/rstb.2014.0169>
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object Perception as Bayesian Inference. *Annu. Rev. Psychol.*, 55, 271–304. <https://doi.org/https://doi.org/10.1146/annurev.psych.55.090902.142005>
- Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology*, 13(2), 150–158. [https://doi.org/10.1016/S0959-4388\(03\)00042-4](https://doi.org/10.1016/S0959-4388(03)00042-4)
- Kilteni, K., & Ehrsson, H. (2020a). Functional connectivity between the cerebellum and somatosensory areas implements the attenuation of self-generated touch. *Journal of Neuroscience*, 40(4), 894–906. <https://doi.org/10.1523/JNEUROSCI.1732-19.2019>
- Kilteni, K., & Ehrsson, H. H. (2017a). Body ownership determines the attenuation of self-generated tactile sensations. *Proceedings of the National Academy of Sciences*, 114(31), 8426–8431. <https://doi.org/10.1073/pnas.1703347114>
- Kilteni, K., & Ehrsson, H. H. (2017b). Sensorimotor predictions and tool use: Hand-held tools attenuate self-touch. *Cognition*, 165, 1–9. <https://doi.org/10.1016/j.cognition.2017.04.005>
- Kilteni, K., & Ehrsson, H. H. (2020b). Predictive attenuation of touch and tactile gating are distinct perceptual phenomena. *BioRxiv*. <https://doi.org/10.1101/2020.11.13.381202>
- Kilteni, K., Houborg, C., & Ehrsson, H. H. (2019). Rapid learning and unlearning of predicted sensory delays in self-generated touch. *ELife*, 8(e42888). <https://doi.org/10.7554/eLife.42888>
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, 83(2). [https://doi.org/10.1016/S0010-0277\(02\)00004-5](https://doi.org/10.1016/S0010-0277(02)00004-5)
- Knill, D., Kersten, D., & Yuille, A. (1996). Introduction: A Bayesian formulation of visual perception. *Perception as Bayesian Inference*, 1, 1–21. <http://diehlab.cps.utexas.edu/Sarah/BayesianTutorial.pdf>
- Kok, P. (2016). Perceptual Inference: A Matter of Predictions and Errors. *Current Biology*, 26(17), R809–R811. <https://doi.org/10.1016/j.cub.2016.07.061>

- Kok, P., Bains, L. J., van Mourik, T., Norris, D. G., & de Lange, F. P. (2016). Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback. *Current Biology*, *26*(3), 371–376. <https://doi.org/10.1016/j.cub.2015.12.038>
- Kok, P., Bains, L. J., Van Mourik, T., Norris, D. G., & De Lange Correspondence, F. P. (2016). Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback. *Current Biology*, *26*(3), 371–376. <https://doi.org/10.1016/j.cub.2015.12.038>
- Kok, P., Brouwer, G. J., Van Gerven, M. A. J., & De Lange, F. P. (2013). Expectations Bias Sensory Representations in Visual Cortex. *Journal of Neuroscience*, *33*(41), 16275–16284. <https://doi.org/10.1523/JNEUROSCI.0742-13.2013>
- Kok, P., & De Lange, F. P. (2015). Predictive Coding in Sensory Cortex. In *An Introduction to Model-Based Cognitive Neuroscience* (pp. 221–244). https://doi.org/10.1007/978-1-4939-2236-9_11
- Kok, P., Failing, M., & de Lange, F. P. (2014). Prior Expectations Evoke Stimulus Templates in the Primary Visual Cortex. *Journal of Cognitive Neuroscience*, *26*(7), 1546–1554.
- Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, *75*(2), 265–270. <https://doi.org/10.1016/j.neuron.2012.04.034>
- Kok, P., Mostert, P., & De Lange, F. P. (2017). Prior expectations induce prestimulus sensory templates. *Proceedings of the National Academy of Sciences*, *114*(39), 10473–10478. <https://doi.org/10.1073/pnas.1705652114>
- Kok, P., Rahnev, D., Jehee, J. F. M., Lau, H. C., & de Lange, F. P. (2012). Attention Reverses the Effect of Prediction in Silencing Sensory Signals. *Cerebral Cortex*, *22*(9), 2197–2206. <https://doi.org/10.1093/cercor/bhr310>
- Kok, P., Rait, L. I., & Turk-Browne, N. B. (2020). Content-based Dissociation of Hippocampal Involvement in Prediction. *Journal of Cognitive Neuroscience*, *32*(3), 527–545. https://doi.org/10.1162/jocn_a_01509
- Kok, P., & Turk-Browne, N. B. (2018). Associative prediction of visual shape in the hippocampus. *Journal of Neuroscience*, *38*(31), 6888–6899. <https://doi.org/10.1523/JNEUROSCI.0163-18.2018>
- Kontaris, I., Wiggett, A. J., & Downing, P. E. (2009). Dissociation of extrastriate body and biological-motion selective areas by manipulation of visual-motor congruency. *Neuropsychologia*, *47*(14), 3118–3124. <https://doi.org/10.1016/j.neuropsychologia.2009.07.012>
- Koopmans, P. J., Barth, M., & Norris, D. G. (2010). Layer-specific BOLD activation in human V1. *Human Brain Mapping*, *31*(9), 1297–1304. <https://doi.org/10.1002/hbm.20936>
- Kusnir, F., Pesin, S., Moscona, G., & Landau, A. N. (2019). When temporal certainty doesn't help. *Journal of Cognitive Neuroscience*, *32*(2), 315–325. https://doi.org/10.1162/jocn_a_01482

- Larsson, J., & Smith, A. T. (2012). fMRI repetition suppression: Neuronal adaptation or stimulus expectation? *Cerebral Cortex*, 22(3), 567–576.
<https://doi.org/10.1093/cercor/bhr119>
- Lawrence, S. J. D., Formisano, E., Muckli, L., & De Lange, F. P. (2019). Laminar fMRI: Applications for cognitive neuroscience. *NeuroImage*, 197, 785–791.
<https://doi.org/10.1016/j.neuroimage.2017.07.004>
- Lawrence, S. J. D., Norris, D. G., & de Lange, F. P. (2019). Dissociable laminar profiles of concurrent bottom-up and top-down modulation in the human visual cortex. *ELife*, 8, e44422. <https://cdn.elifesciences.org/articles/44422/elifesciences/44422-v2.pdf>
- Lawrence, S. J. D., van Mourik, T., Kok, P., Koopmans, P. J., Norris, D. G., & de Lange, F. P. (2018). Laminar Organization of Working Memory Signals in Human Visual Cortex. *Current Biology*, 28(21), 3435–3440.e4.
<https://doi.org/10.1016/J.CUB.2018.08.043>
- Lee, T. S., & Nguyen, M. (2001). Dynamics of subjective contour formation in the early visual cortex. *Proceedings of the National Academy of Sciences*, 98(4), 1907–1911. <https://doi.org/10.1073/pnas.031579998>
- Leite, F. P., & Ratcliff, R. (2011). What cognitive processes drive response biases? A diffusion model analysis. *Judgment and Decision Making*, 6(7), 651–687.
<https://psycnet.apa.org/record/2011-24967-007>
- Lewis, R., & Noppeney, U. (2010). Audiovisual Synchrony Improves Motion Discrimination via Enhanced Connectivity between Early Visual and Auditory Areas. *Journal of Neuroscience*, 30(37), 12329–12339.
<https://doi.org/10.1523/JNEUROSCI.5745-09.2010>
- Livesey, E. J., & McLaren, I. P. (2017). Revisiting peak shift on an artificial dimension: Effects of stimulus variability on generalisation. *Quarterly Journal of Experimental Psychology*, 72(2), 132–150.
<https://doi.org/10.1177/1747021817739832>
- Mackintosh, N. J., & Little, L. (1969). Intradimensional and extradimensional shift learning by pigeons. *Psychonomic Science*, 14(1), 5–6.
- Markuerkiaga, I., Barth, M., & Norris, D. G. (2016). A cortical vascular model for examining the specificity of the laminar BOLD signal. *NeuroImage*, 132, 491–498.
<https://doi.org/10.1016/j.neuroimage.2016.02.073>
- McDonough, K. L., Hudson, M., & Bach, P. (2019). Cues to intention bias action perception toward the most efficient trajectory. *Scientific Reports*, 9(1).
<https://doi.org/10.1038/s41598-019-42204-y>
- Mechelli, A., Price, C. J., Noppeney, U., & Friston, K. J. (2003). A Dynamic Causal Modeling Study on Category Effects: Bottom-Up or Top-Down Mediation? *Journal of Cognitive Neuroscience*, 15(7), 925–934.
<https://doi.org/10.1162/089892903770007317>
- Moerel, M., De Martino, F., Kemper, V. G., Schmitter, S., Vu, A. T., Uğurbil, K., Formisano, E., & Yacoub, E. (2018). Sensitivity and specificity considerations for fMRI encoding, decoding, and mapping of auditory cortex at ultra-high field HHS

- Public Access. *Neuroimage*, *164*, 18–31.
<https://doi.org/10.1016/j.neuroimage.2017.03.063>
- Muckli, L., De Martino, F., Ugurbil, K., Goebel, R., & Essa Yacoub. (2015). Contextual Feedback to Superficial Layers of V1. *Current Biology*, *25*, 2690–2695.
<https://doi.org/10.1016/j.cub.2015.08.057>
- Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff. *Journal of Neuroscience*, *32*(7), 2335–2343.
<https://doi.org/10.1523/JNEUROSCI.4156-11.2012>
- Noppeney, U., Josephs, O., Hocking, J., Price, C. J., & Friston, K. J. (2008). The Effect of Prior Visual Information on Recognition of Speech and Sounds. *Cerebral Cortex*, *18*(3), 598–609. <https://doi.org/10.1093/cercor/bhm091>
- Ofen, N., Moran, A., & Sagi, D. (2007). Effects of trial repetition in texture discrimination. *Vision Research*, *47*(8), 1094–1102.
<https://doi.org/10.1016/j.visres.2007.01.023>
- Palmer, T. E. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition*, *3*(5), 519–526. <https://doi.org/10.3758/BF03197524>
- Paraskevoudi, N., & SanMiguel, I. (2020). Self-generation and sound intensity interactively modulate perceptual bias, but not perceptual sensitivity. *BioRxiv*.
<https://doi.org/10.1101/2020.11.23.393785>
- Pinto, Y., van Gaal, S., de Lange, F. P., Lamme, V. A. F., & Seth, A. K. (2015). Expectations accelerate entry of visual stimuli into awareness. *Journal of Vision*, *15*(8), 13–13. <https://doi.org/10.1167/15.8.13>
- Posner, M., Snyder, C., & Davidson, B. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, *109*(2), 160–174.
<https://psycnet.apa.org/record/1981-11809-001>
- Press, C., & Cook, R. (2015). Beyond action-specific simulation: domain-general motor contributions to perception. *Trends in Cognitive Sciences*, *19*(4), 176–178.
<https://doi.org/10.1016/j.tics.2015.01.006>
- Press, C., Gherri, E., Heyes, C., & Eimer, M. (2010). Action Preparation Helps and Hinders Perception of Action. *Journal of Cognitive Neuroscience*, *22*(10), 2198–2211.
- Press, C., Kok, P., & Yon, D. (2020a). Learning to Perceive and Perceiving to Learn. *Trends in Cognitive Sciences*, *24*(4), 260–261.
<https://doi.org/10.1016/j.tics.2020.01.002>
- Press, C., Kok, P., & Yon, D. (2020b). The Perceptual Prediction Paradox. *Trends in Cognitive Sciences*, *24*(1), 13–24. <https://doi.org/10.1016/J.TICS.2019.11.003>
- Press, C., & Yon, D. (2019). Perceptual Prediction: Rapidly Making Sense of a Noisy World. *Current Biology*, *29*(15), R751–R753.
<https://doi.org/10.1016/j.cub.2019.06.054>
- Prins, N., & Kingdom, F. A. A. (2018). Applying the model-comparison approach to

- test specific research hypotheses in psychophysical research using the Palamedes toolbox. *Frontiers in Psychology*, 9(1250).
<https://doi.org/10.3389/fpsyg.2018.01250>
- Puri, A. M., & Wojciulik, E. (2008). Expectation both helps and hinders object perception. *Vision Research*, 48(4), 589–597.
<https://doi.org/10.1016/j.visres.2007.11.017>
- Rahnev, D., Lau, H., & De Lange, F. P. (2011). Prior Expectation Modulates the Interaction between Sensory and Prefrontal Regions in the Human Brain. *The Journal of Neuroscience*, 31(29), 10741–10748.
<https://doi.org/10.1523/JNEUROSCI.1478-11.2011>
- Ramírez, F. M., & Merriam, E. P. (2020). Forward models of repetition suppression depend critically on assumptions of noise and granularity. *Nature Communications*, 11(1), 4732. <https://doi.org/10.1038/s41467-020-18315-w>
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87. <https://doi.org/10.1038/4580>
- Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, 20(4), 873–922.
<https://doi.org/10.1162/neco.2008.12-06-420>
- Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations on the effectiveness of reinforcement and non-reinforcement. In *Classical conditioning II: Current research and theory*.
- Reznik, D., Henkin, Y., Levy, O., & Mukamel, R. (2015). Perceived Loudness of Self-Generated Sounds Is Differentially Modified by Expected Sound Intensity. *PLOS ONE*, 10(5), e0127651. <https://doi.org/10.1371/journal.pone.0127651>
- Reznik, D., Henkin, Y., Schadel, N., & Mukamel, R. (2014). Lateralized enhancement of auditory cortex activity and increased sensitivity to self-generated sounds. *Nature Communications*, 5(1), 4059. <https://doi.org/10.1038/ncomms5059>
- Richter, D., & de Lange, F. P. (2019). Statistical learning attenuates visual activity only for attended stimuli. *ELife*, 8, e47869. <https://doi.org/10.1101/653782>
- Richter, D., Ekman, M., & de Lange, F. P. (2018). Suppressed sensory response to predictable object stimuli throughout the ventral visual stream. *Journal of Neuroscience*, 38(34), 7452–7461. <https://doi.org/10.1523/JNEUROSCI.3421-17.2018>
- Rockland, K. S., & Virga, A. (1989). Terminal arbors of individual “Feedback” axons projecting from area V2 to V1 in the macaque monkey: A study using immunohistochemistry of anterogradely transported Phaseolus vulgaris-leucoagglutinin. *The Journal of Comparative Neurology*, 285(1), 54–72.
<https://doi.org/10.1002/cne.902850106>
- Rohenkohl, G., Gould, I. C., Pessoa, J., & Nobre, A. C. (2014). Combining spatial and temporal expectations to improve visual perception. *Journal of Vision*, 14(4), 8–8.
<https://doi.org/10.1167/14.4.8>

- Roussel, C., Hughes, G., & Waszak, F. (2014). Action prediction modulates both neurophysiological and psychophysical indices of sensory attenuation. *Frontiers in Human Neuroscience*, 8, 115. <https://doi.org/10.3389/fnhum.2014.00115>
- Schapiro, A., & Turk-Browne, N. (2015). Statistical Learning. *Brain Mapping: An Encyclopedic Reference*, 3, 501–506. <https://doi.org/10.1016/B978-0-12-397025-1.00276-1>
- Schneider, D. M., Sundararajan, J., & Mooney, R. (2018). A cortical filter that learns to suppress the acoustic consequences of movement. *Nature*, 561, 391–395. <https://doi.org/10.1038/s41586-018-0520-5>
- Schwarz, K. A., Pfister, R., Kluge, M., Weller, L., & Kunde, W. (2017). Do We See It or Not? Sensory Attenuation in the Visual Domain. *Journal of Experimental Psychology: General*, 147(3), 418–430. <https://doi.org/10.1037/xge0000353>
- Seki, K., & Fetz, E. E. (2012). Gating of Sensory Input at Spinal and Cortical Levels during Preparation and Execution of Voluntary Movement. *Journal of Neuroscience*, 32(3), 890–902. <https://doi.org/10.1523/JNEUROSCI.4958-11.2012>
- Self, M. W., van Kerkoerle, T., Supèr, H., & Roelfsema, P. R. (2013). Distinct roles of the cortical layers of area V1 in figure-ground segregation. *Current Biology*, 23(21), 2121–2129. <https://doi.org/10.1016/j.cub.2013.09.013>
- Shanks, D. R., & Dickinson, A. (1991). Instrumental judgment and performance under variations in action-outcome contingency and contiguity. *Memory & Cognition*, 19(4), 353–360. <https://doi.org/10.3758/BF03197139>
- Shea, N. (2015). Distinguishing Top-Down From Bottom-Up Effects. In *Perception & Its Modalities*.
- Shergill, S. S., Bays, P. M., Frith, C. D., & Wolpert, D. M. (2003). Two Eyes for an Eye: The Neuroscience of Force Escalation. *Science*, 301(5630), 187. <https://doi.org/10.1126/science.1085327>
- Shergill, S. S., Samson, G., Bays, P. M., Frith, C. D., & Wolpert, D. M. (2005). Evidence for sensory prediction deficits in schizophrenia. *American Journal of Psychiatry*, 162(12), 2384–2386. <https://doi.org/10.1176/appi.ajp.162.12.2384>
- Shergill, S. S., White, T. P., Joyce, D. W., Bays, P. M., Wolpert, D. M., & Frith, C. D. (2013). Modulation of somatosensory processing by action. *NeuroImage*, 70, 356–362. <https://doi.org/10.1016/J.NEUROIMAGE.2012.12.043>
- Shergill, S. S., White, T. P., Joyce, D. W., Bays, P. M., Wolpert, D. M., & Frith, C. D. (2014). Functional magnetic resonance imaging of impaired sensory prediction in schizophrenia. *JAMA Psychiatry*, 71(1), 28–35. <https://doi.org/10.1001/jamapsychiatry.2013.2974>
- Singla, S., Dempsey, C., Warren, R., Enikolopov, A. G., & Sawtell, N. B. (2017). A cerebellum-like circuit in the auditory system cancels responses to self-generated sounds. *Nature Neuroscience*, 20(7), 943–950. <https://doi.org/10.1038/nn.4567>
- Smith, F. W., & Muckli, L. (2010). Nonstimulated early visual areas carry information about surrounding context. *Proceedings of the National Academy of Sciences of the United States of America*, 107(46), 20099–20103.

<https://doi.org/10.1073/pnas.1000233107>

- Stanley, J., & Miall, R. C. (2007). Functional activation in parieto-premotor and visual areas dependent on congruency between hand movement and visual stimuli during motor-visual priming. *NeuroImage*, *34*(1), 290–299. <https://doi.org/10.1016/j.neuroimage.2006.08.043>
- Stephan, K. E., Petzschner, F. H., Kasper, L., Bayer, J., Wellstein, K. V., Stefanics, G., Pruessmann, K. P., & Heinzle, J. (2019). Laminar fMRI and computational theories of brain function. *NeuroImage*, *197*, 699–706. <https://doi.org/10.1016/j.neuroimage.2017.11.001>
- Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), 745–756. <https://doi.org/10.1038/nrn3838>
- Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*, *13*(9), 403–409. <https://doi.org/10.1016/j.tics.2009.06.003>
- Summerfield, C., & Egner, T. (2016). Feature-Based Attention and Feature-Based Expectation. *Trends in Cognitive Sciences*, *20*(6), 401–404. <https://doi.org/10.1016/j.tics.2016.04.003>
- Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M.-M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, *11*(9), 1004. <https://doi.org/10.1038/nn.2163>
- Todorovic, A., & De Lange, F. P. (2012). Repetition Suppression and Expectation Suppression Are Dissociable in Time in Early Auditory Evoked Fields. *The Journal of Neuroscience*, *32*(39), 13389–13395. <https://doi.org/10.1523/JNEUROSCI.2227-12.2012>
- Tukey, J. W. (1962). The Future of Data Analysis. *The Annals of Mathematical Statistics*, *33*(1), 1–67.
- Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit Perceptual Anticipation Triggered by Statistical Learning. *Journal of Neuroscience*, *30*(33), 11177–11187. <https://doi.org/10.1523/JNEUROSCI.0858-10.2010>
- Urai, A. E., de Gee, J. W., Tsetsos, K., & Donner, T. H. (2019). Choice history biases subsequent evidence accumulation. *eLife*, *8*(e46331). <https://doi.org/10.7554/eLife.46331>
- van Ede, F., van Doren, T. I., Damhuis, J., de Lange, F. P., & Maris, E. (2015). Movement preparation improves touch perception without awareness. *Cognition*, *137*, 189–195. <https://doi.org/10.1016/j.cognition.2015.01.009>
- van Kerkoerle, T., Self, M. W., & Roelfsema, P. R. (2017). Layer-specificity in the effects of attention and working memory on activity in primary visual cortex. *Nature Communications*, *8*(1), 1–14. <https://doi.org/10.1038/ncomms13804>
- van Mourik, T., Koopmans, P. J., & Norris, D. G. (2019). Improved cortical boundary registration for locally distorted fMRI scans. *PLOS ONE*, *14*(11), e0223440.

<https://doi.org/10.1371/journal.pone.0223440>

- van Mourik, T., van der Eerden, J. P. J. M., Bazin, P.-L., & Norris, D. G. (2019). Laminar signal extraction over extended cortical areas by means of a spatial GLM. *PLOS ONE*, *14*(3), e0212493. <https://doi.org/10.1371/journal.pone.0212493>
- von Helmholtz, H. (1867). *Handbuch der physiologischen Optik*. Leipzig, Voss.
- Voudouris, D., Broda, M. D., & Fiehler, K. (2019). Anticipatory grasping control modulates somatosensory perception. *Journal of Vision*, *19*(5), 1–10. <https://doi.org/10.1167/19.5.4>
- Wagstyl, K., Larocque, S., Cucurull, G., Lepage, C., Cohenid, J. P., Bludaid, S., Palomero-Gallagherid, N., Lewisid, L. B., Funck, T., Spitzerid, H., Dickscheidid, T., Fletcher, P. C., Romeroid, A., Zillesid, K., Amunts, K., Bengioid, Y., & Evansid, A. C. (2020). BigBrain 3D atlas of cortical layers: Cortical and laminar thickness gradients diverge in sensory and motor cortices. *PLoS Biology*, *18*(4), e3000678. <https://doi.org/10.1371/journal.pbio.3000678>
- Webster, M. A. (2015). Visual Adaptation. *Annual Review of Vision Science*, *1*(1), 547–567. <https://doi.org/10.1146/annurev-vision-082114-035509>
- Weiss, C., Herwig, A., & Schütz-Bosbach, S. (2011). The self in action effects: Selective attenuation of self-generated sounds. *Cognition*, *121*(2), 207–218. <https://doi.org/10.1016/j.cognition.2011.06.011>
- Whitford, T. J., Jack, B. N., Pearson, D., Griffiths, O., Luque, D., Harris, A. W., Spencer, K. M., & Le Pelley, M. E. (2017). Neurophysiological evidence of efference copies to inner speech. *ELife*, *6*, e28197. <https://doi.org/10.7554/eLife.28197>
- Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Frontiers in Neuroinformatics*, *7*, 14. <https://doi.org/10.3389/fninf.2013.00014>
- Williams, S. R., & Chapman, E. C. (2000). Time-course and magnitude of movement-related gating of tactile detection in humans. II. Effects of Stimulus Intensity. *Journal of Neurophysiology*, *84*(2), 863–875. <https://doi.org/https://doi.org/10.1152/jn.2002.88.4.1968>
- Williams, S. R., Shenasa, J., & Chapman, E. C. (1998). Time course and magnitude of movement-related gating of tactile detection in humans. I. Importance of stimulus location. *Journal of Neurophysiology*, *79*(2), 947–963. <https://doi.org/10.1152/jn.00527.2001>
- Wills, S., & Mackintosh, N. J. (1998). Peak Shift on an Artificial Dimension. *The Quarterly Journal of Experimental Psychology Section B*, *51*(1), 1–32. <https://doi.org/10.1080/713932669>
- Wolpe, N., Ingram, J. N., Tsvetanov, K. A., Geerligs, L., Kievit, R. A., Henson, R. N., Wolpert, D. M., Rowe, J. B., Tyler, L. K., Brayne, C., Bullmore, E., Calder, A., Cusack, R., Dalgleish, T., Duncan, J., Matthews, F. E., Marslen-Wilson, W., Shafto, M. A., Campbell, K., ... Rowe, J. B. (2016). Ageing increases reliance on sensorimotor prediction through structural and functional differences in

- frontostriatal circuits. *Nature Communications*, 7(1), 13034.
<https://doi.org/10.1038/ncomms13034>
- Wolpe, N., Zhang, J., Nombela, C., Ingram, J. N., Wolpert, D. M., & Rowe, J. B. (2018). Sensory attenuation in Parkinson's disease is related to disease severity and dopamine dose. *Scientific Reports*, 8(1), 15643. <https://doi.org/10.1038/s41598-018-33678-3>
- Wolpert, D. M., & Flanagan, J. R. (2001). Motor prediction. *Current Biology*, 11(18), R729–R732. [https://doi.org/10.1016/S0960-9822\(01\)00432-8](https://doi.org/10.1016/S0960-9822(01)00432-8)
- Wyart, V., Nobre, A. C., & Summerfield, C. (2012). Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proceedings of the National Academy of Sciences*, 109(16), 3593–3598.
<https://doi.org/10.1073/pnas.1204601109>
- Yon, D., Gilbert, S. J., De Lange, F. P., & Press, C. (2018). Action sharpens sensory representations of expected outcomes. *Nature Communications*, 9(1), 4288.
<https://doi.org/10.1038/s41467-018-06752-7>
- Yon, D., & Press, C. (2017). Predicted action consequences are perceptually facilitated before cancellation. *Journal of Experimental Psychology: Human Perception and Performance*, 43(6), 1073–1083. <https://doi.org/10.1037/xhp0000385>
- Yon, D., & Press, C. (2018). Sensory predictions during action support perception of imitative reactions across suprasecond delays. *Cognition*, 173, 21–27.
<https://doi.org/doi:10.1016/j.cognition.2017.12.008>
- Yon, D., Zainzinger, V., Lange, F. de, Eimer, M., & Press, C. (2020). Action biases perceptual decisions toward expected outcomes. *Journal of Experimental Psychology: General*, Advance online publication.
<https://doi.org/10.1037/xge0000826>
- Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences*, 10(7), 301–308.
<https://doi.org/10.1016/j.tics.2006.05.002>
- Zamboni, E., Ledgeway, T., McGraw, P. V., & Schluppeck, D. (2016). Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proceedings of the Royal Society B: Biological Sciences*, 283(1833), 20160263. <https://doi.org/10.1098/rspb.2016.0263>