

PERCEPTUAL LEARNING FOR EXPERTISE

by

MADELEINE RANSOM

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

In

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES

(Philosophy)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

March 2020

© Madeleine Ransom, 2020

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation entitled:

Perceptual Learning for Expertise

submitted by Madeleine Ransom in partial fulfillment of the requirements for

the degree of Doctor of Philosophy

in Philosophy

Examining Committee:

Dominic McIver Lopes, Professor, Department of Philosophy, UBC

Co-supervisor

Murat Aydede, Professor, Department of Philosophy, UBC

Co-supervisor

Christopher Mole, Professor, Department of Philosophy, UBC

Supervisory Committee Member

Eric Margolis, Professor, Department of Philosophy, UBC

University Examiner

Lawrence Ward, Professor, Department of Psychology, UBC

University Examiner

Additional Supervisory Committee Members:

Evan Thompson, Professor, Department of Philosophy, UBC

Supervisory Committee Member

Abstract

What does the expert have that the novice does not? One component of expertise may be perceptual, involving a change in what we are able to perceive. Experts develop the ability to taste the subtle flavours of a wine, hear minute variations in pitch imperceptible to novices, or distinguish shades of red that are indistinguishable to novices. More controversially, experts perceive a bird to be a northern flicker, a shadow on an x-ray to be a tumour, or a painting to be beautiful.

This is controversial because a competing explanation is that experts merely apply their extensive background knowledge to selectively attend to the relevant aspects of their perceptual experience, and then make such judgments in cognition. On this explanation there is no substantive change in perceptual experience between novice and expert.

In this dissertation I argue against this alternate explanation of expert ability and defend the perceptual expertise thesis: through perceptual learning experts come to perceive high-level properties imperceptible to novices. I do this in part by appealing to empirical studies of perceptual learning and perceptual expertise. The positive account of perceptual expertise I build here allows for the resolution of a puzzle in aesthetics regarding the role of training, and a clarification of the epistemic significance of perceptual learning.

Lay Summary

Can practice make us better at perceiving? The scientific study of perceptual learning suggests the answer to the question is yes. However, it is not clear exactly what this learning consists in. My dissertation clarifies how our perceptual experiences change with expertise. I argue that we can get better at perceiving things like shapes and colours, and that we can also come to directly perceive – rather than infer – things like F-35 fighter planes, tumours in x-rays, and the beauty of artworks. This advances our understanding of the nature of expertise, and so can inform how training programs are designed in fields from the military to medicine to art criticism.

Preface

The work contained in this dissertation is original to me, and was not produced with any co-authors. I am responsible for the identification and design of my research program, the performance of all research contained in this document, and for the analysis of the research data.

Chapter four -“Learning to See Beauty” contains some themes and ideas that also appear in a published paper, with the reference as follows: Ransom, Madeleine. 2020. “Waltonian Perceptualism.” *Journal of Aesthetics and Art Criticism* 78 (1): 66–70.

Table of Contents

Abstract.....	iii
Lay Summary	iv
Preface.....	v
Table of Contents	vi
Acknowledgements	ix
Dedication	xi
Chapter 1: Introduction	1
1.1 The content debate	2
1.2 The method of phenomenal contrast.....	3
1.3 Perceptual learning.....	5
1.4 Perceptual expertise	6
1.5 Structure of this dissertation	8
1.5.1 Using empirical methods to adjudicate the content debate.....	8
1.5.2 What is learned in perceptual learning.....	9
1.5.3 Aesthetic perception.....	11
1.6 The epistemology of perceptual learning.....	12
1.7 Significance of this work	14
1.7.1 Real-life expertise training.....	14
1.7.2 Visual literacy	15
1.7.3 The perception-cognition divide.....	16
Chapter 2: Using Empirical Methods to Determine the Contents of Perception	18
2.1 Introduction.....	19
2.2 The common mechanism argument.....	20
2.2.1 Adaptation criterion	22
2.2.2 Automaticity criterion.....	22
2.2.3 Encapsulation criterion	23
2.2.4 Attention criterion.....	24
2.2.5 Speed criterion	26
2.3 Scepticism about the prospects of empirical methodology.....	29
2.4 The converse fallacy pitfall.....	32
2.5 Low-level confound pitfall	37
2.6 Interpreting the results	41
2.6.1 Asymmetry.....	41
2.6.2 Generalizing to a class	44
2.6.3 Conflicting results.....	46
2.7 Conclusion	48
Chapter 3: What is Learned in Perceptual Learning.....	51
3.1 Introduction.....	52
3.1.1 Why does the debate matter?	53
3.1.2 What is perceptual learning?.....	55
3.2 Perceptual categories are high-level	58

3.2.1	The post-perceptual judgment hypothesis	59
3.2.2	What counts as a high-level property.....	61
3.2.3	Argument from the mechanics of perceptual learning.....	62
3.3	The learned properties are represented in perception	72
3.3.1	Speed.....	75
3.3.2	Automaticity and encapsulation.....	77
3.3.3	Adaptation effects	79
3.4	There is a change in the phenomenal character of perceptual experience due to high-level properties.....	81
3.5	Conclusion	87
Chapter 4: Learning to See Beauty: A novel account of aesthetic perception		89
4.1	Introduction.....	90
4.2	Considerations for and against aesthetic perception.....	92
4.2.1	Language.....	92
4.2.2	Psychological immediacy	92
4.2.3	Aesthetic diversity	93
4.2.4	Non-perceptual art forms and objects of aesthetic attribution.....	94
4.2.5	The evaluative nature of aesthetic properties.....	95
4.2.6	Culturally diverse and changing aesthetic ideals	96
4.2.7	Aesthetic training.....	97
4.2.8	Lack of aesthetic training.....	99
4.3	Arguments in favour of aesthetic perception	101
4.3.1	Phenomenal contrast argument.....	101
4.3.2	Can we appeal to empirical data?	103
4.3.3	How to move forward: providing a positive account.....	104
4.4	Aesthetic perception by cognitive permeation.....	105
4.4.1	Walton's categories of art.....	107
4.4.2	'Categories of Art' as cognitive permeation.....	108
4.4.3	Against 'Categories of Art' as cognitive permeation	110
4.5	Waltonian perceptualism	114
4.5.1	Perceptual learning.....	115
4.5.2	Fluency hypothesis.....	121
4.5.3	Perceptual expertise	125
4.6	Meeting the desiderata	128
4.6.1	Cultural desideratum.....	128
4.6.2	Diversity desideratum	129
4.6.3	Abstracta desideratum.....	129
4.6.4	Affect desideratum.....	131
4.6.5	Training desideratum	132
4.6.6	Un-trained desideratum.....	133
4.7	Outstanding issues	134
4.7.1	What is the correct category for viewing objects in?.....	134
4.7.2	Are aesthetic judgments that do not rely on categories still possible?	135
4.7.3	Is there a limit to which categories we can become perceptual experts in?	135
4.7.4	The naked mole rat problem	136

4.7.5	The Cindy Crawford mole problem.....	137
4.7.6	How can this account be extended to other aesthetic properties?.....	138
4.8	Conclusion	140
Chapter 5: The Epistemology of Perceptual Learning.....		141
5.1	Introduction.....	142
5.1.1	The content debate	143
5.1.2	Immediate justification	145
5.2	Cognitive permeation.....	147
5.2.1	Defining cognitive permeation	148
5.2.2	Diachronic cognitive permeation.....	150
5.2.2.1	Diachronic cognitive permeation via attention	153
5.2.3	Epistemic implications of cognitive permeation.....	156
5.3	Perceptual expertise is not the result of cognitive permeation	160
5.3.1	Perceptual learning.....	161
5.3.2	Arguing for the pattern-driven hypothesis.....	164
5.3.3	The role of cognition in perceptual learning.....	175
5.4	Categorization vs. concepts	178
5.5	Conclusion	182
Chapter 6: Conclusion.....		185
6.1	Using empirical methods to resolve the content debate.....	185
6.2	What is learned in perceptual learning.....	186
6.3	Learning to see beauty: a novel account of aesthetic perception.....	188
6.4	The epistemology of perceptual learning.....	188
6.5	Directions for future research	189
6.5.1	Bias in perceptual learning.....	190
6.5.2	Perceptual learning and theory-ladenness.....	195
References.....		202

Acknowledgements

During my time at UBC, I feel very lucky to have worked with people who are not only great philosophers, but also great human beings. First I would like to thank my supervisory committee. My utmost gratitude to Dominic McIver Lopes, without whose unwavering support, guidance and wisdom I could not have completed this project. Dom is pretty much the Platonic ideal form of a supervisor. Thanks to Murat Aydede for our many café conversations that always provided me with a new, more nuanced perspective on the issue (and a long reading list to follow up on). Thanks to Christopher Mole for challenging me on so many occasions to make my arguments stronger and my positions clearer. I am also grateful to him for his collaboration on a paper on attention and predictive processing that taught me about co-authorship and the publishing process. And finally, thanks to Evan Thompson for our many interesting discussions on attention and for his collaboration on a project on affect and predictive processing.

Within the Philosophy Department I would like to extend my gratitude in particular to Eric Margolis for mentoring me during a one and a half year teaching certificate program, and also for all his help and career advice. He went above and beyond as placement director in making sure students were prepared for the job market. I would also like to thank Chris Stephens for his informal mentorship, exemplary teaching, and feedback in my development as a teacher. Finally, my gratitude to John Woods and Adam Morton, both of whom always took the time to attend my work-in-progress seminars and give me extensive feedback on my papers.

Outside the department, I would like to thank David Silver at the W. Maurice Young Centre for Applied Ethics for his mentorship and support during my fellowship there. Conversations with him, other faculty, and the other fellows of the Centre helped broaden my

view of the role of ethics, and have greatly influenced the course of my post-dissertation research projects. I would also like to thank Beck Todd in the Department of Psychology for allowing me to be part of her lab for over a year. I learned much about psychology and neuroscience during my time there. I owe special thanks to my undergraduate professors, in particular David Albert and Derek Matravers, who both encouraged me to go into philosophy. I would also like to thank my grade four teacher, Judy Kyle for introducing me to categorical logic. This spawned my interest in philosophy, and impacted my decision to take my first philosophy class in university. My time at UBC was greatly enriched by the grad students of the philosophy department. I owe particular thanks to Sina Fazelpour, Servaas van der Berg, Aleksey Balotskiy, Emma Esmaili, Kousaku Yui, Richard Sandlin, Jasper Heaton, Jelena Markovic, and Matthew Smithdeal for their friendship and support.

I of course would not be where I am today without the support of my family. Thank you all for helping me to become the person I am. In particular thanks to Bruce, Keira, and Blaise for putting up with all my long hours, and to my grandparents, parents, in-laws, brother and sisters, cousins and second cousins, nieces and nephews, aunts and uncles for your friendship. A special thanks also to my great uncle Władysław Tatarkiewicz for being a source of inspiration (and making my career choice seem more reasonable to my parents). Thank you to my friends outside academia, and in particular Claudia Tiscareño, Karly Curcio Gula, Katherine McMillin, and Ione Doig Smith.

Finally, I would not have been able to complete this dissertation without financial support from a SSHRC Bombardier scholarship, the American Society for Aesthetics dissertation fellowship, and a fellowship at the W. Maurice Young Centre for Applied Ethics. My utmost gratitude to these institutions for supporting this research.

To my parents, for giving me the freedom in life to discover my own path and for supporting me along the way. I love you both so much.

Chapter 1: Introduction

I defend the perceptual expertise thesis: through perceptual learning experts come to perceive high-level properties imperceptible to novices.

Experts in diverse fields appear to be able to perceive things that novices cannot. Sommeliers report tasting properties of a wine that casual drinkers are blind to. Whereas novices may describe a wine as simply tasting sweet, tangy, or sour, experts may describe it as tasting oaky, balanced, or with notes of citrus. Trained musicians can distinguish pitches that novices are at chance in distinguishing, and can also discern patterns that novices appear unable to. Expert ornithologists can distinguish between birds that ‘look the same’ to novice birdwatchers.

Reflecting on such cases of expertise raises two central questions that will serve to guide this project. First, is there really a difference in the perceptual experience of experts as the result of training, or have experts simply learned to do more with the perceptual information they receive, perhaps employing background knowledge to make more sophisticated inferences? Second, supposing that experts do have a different perceptual experience from novices, what sorts of properties make the difference? Do they merely perceive a wider array of simple perceptual properties such as colour, motion, shape, and size, or do they also perceive more complex properties such as the property of being a pileated woodpecker, a balanced wine, a fugue, or a Picasso?

These two guiding questions are informed and rendered more precise by the content debate in the philosophy of perception.

1.1 The content debate

To say that perceptual experiences have representational contents is to say that such experiences can be assessed for accuracy – they convey to the perceiver that the world is a certain way, where what is conveyed can either be an accurate or inaccurate representation of the world. One way in which these contents can be assessed for accuracy is in terms of the properties they ascribe – truly or falsely – to objects. The content debate concerns what sorts of properties can be part of perceptual experience, where properties are separated into two classes: low and high-level. Low-level properties are usually defined stipulatively – they are the properties that everyone agrees are represented in perceptual experience, and include shapes, colours, motion, relative position, and the like. High-level properties are all those properties that are not low-level. They include natural and artifactual kind properties such as being a tiger or a teacup, aesthetic and moral properties such as being beautiful or good, and properties such as being caused by a given event, amongst others. It is controversial whether high-level properties can be part of the contents of perceptual experience.

Those who hold the position that we perceptually experience only low-level properties are sparse content theorists. They hold that we perceptually experience only low-level properties, such as the ones listed above, along with analogues in the other sensory modalities (Tye 1995; 2000; Dretske 1981; Clark 2000; Raftopoulos 2009; Connolly 2014; 2019). Those who hold the position that we perceptually experience at least some high-level properties are rich content

theorists (Siegel 2010; 2006b; Fish 2013; Werner 2016; Siewert 1998; Brogaard 2018; Stokes 2018a).¹

Our two questions regarding expertise can now be rephrased in terms of the content of perceptual experience. First, can training or experience alter the contents of perceptual experience? Second, if it can alter the contents of perceptual experience, is it with respect to low-level properties only, or also with respect to high-level properties? An initial, prominent attempt to answer these questions, and one that serves as a point of departure for my own work, is the method of phenomenal contrast.

1.2 The method of phenomenal contrast

The most discussed defence of the rich content view – embedded in the work that is chiefly responsible for illuminating the importance of the debate – is the method of phenomenal contrast (Siegel 2006b; 2010; see also Siewert 1998). The argument rests on the intuition that there is a difference in phenomenology between the experience of the novice and expert, or before and after training of some sort. For example, take someone who is about to learn to speak Farsi. It is plausible that after she has achieved some level of proficiency that the phenomenology of hearing spoken Farsi is different from her experience prior to learning. Or, take a novice and an expert birder observing a northern flicker out in the field. Again, it is

¹ Here I set aside skeptical worries that there is no way of settling the issue (Lycan 2014) or that there is no fact of the matter as to whether we

plausible that there is a phenomenological difference between them, one that has something to do with the expert's having learned to recognize northern flickers.

Granting this intuition, Siegel (2010) labels the overall experience of a subject S prior to training 'the contrasting experience' and the overall experience of S after training 'the target experience.' E1 is the visual experience had by S prior to training, and E2 is the visual experience had by S after training. E1 will thus be part of the contrasting experience, and E2 part of the target experience. Her argument then proceeds along the following lines (Siegel 2010, 101):

- (1) The target experience differs in its phenomenology from the contrasting experience.
- (2) If the target experience differs in its phenomenology from the contrasting experience, then there is a phenomenological difference between E1 and E2.
- (3) If there is a phenomenological difference between E1 and E2, then E1 and E2 differ in content.
- (4) If there is a difference in content between E1 and E2, it is a difference with respect to [high-level] properties represented in E1 and E2.
- (5) Therefore, there is a difference with respect to high-level properties represented in E1 and E2.

However, sparse content theorists are unlikely to grant the "minimal intuition one has to have for the argument to get off the ground" (Siegel, 2010, p. 101). That is, they may deny that there is a difference in what it is like to see the relevant object upon developing the expert recognitional capacity (Dretske 2015; Lyons 2005). The change between novice and expert might be said instead to concern what sorts of inferences or associations the expert makes –

while perceptual experience remains the same, other things change that explain the difference. Though Siegel has provided arguments against other such possibilities (2010, p. 101-113), there has also been much pushback (Butterfill 2009; Pautz 2009; Crutchfield 2012; Reiland 2014; Furst 2017; Price 2009; O'Callaghan 2011; Koksvik 2015; Brogaard 2013). Such a persistent clash of intuitions may even be reason to mistrust the argument entirely, if the intuitions are being produced by one's background theory (Fish 2013).

In the face of this stalemate, I draw upon the real-world psychological study of perceptual learning and perceptual expertise to advance the content debate.

1.3 Perceptual learning

Perceptual learning provides the resources to answer our first guiding question positively: there can be a change in perceptual experience between novice and expert. While it was long thought that the perceptual system remained relatively fixed after a short 'critical period' in early development (Wiesel and Hubel 1963; Hubel and Wiesel 1970), there is a wide body of evidence that attests to its plasticity well into adulthood (Hooks and Chen 2007). One branch of study of this plasticity is perceptual learning. Following Gibson (1963) and Goldstone (1998), perceptual learning involves structural and functional changes in the perceptual system, due to repeated exposure to a stimulus, that result in a change in perceptual experience (see also Connolly 2019; Kellman and Garrigan 2009). Importantly, perceptual learning is perceptual: the changes brought about must make a difference to how something is perceived, or whether it is perceived at all. Moreover, the changes must be brought about as a result of learning, which discounts changes

due to lesions and aging, for example. Learning must also take place as a result of repeated exposure to or practice with the target stimuli.

Some perceptual learning can be wholly explained as changes in the ability to perceive low-level properties. For example, the ability to discriminate the direction of motion of dots on a screen improves with extensive perceptual practice with the task (Ball and Sekuler 1987; 1982; Matthews et al. 1999). In such cases, the improvement is uncontroversially perceptual – experimental design ensures that subjects are not getting better just by guessing or inferring direction of motion.

However, in such cases improvements in subjects' abilities are not explained by the representation of any high-level properties in perceptual experience. To answer the second guiding question, we must go beyond cases of 'low-level' perceptual learning. There is also empirical work that suggests that we can come to represent high-level properties in perceptual experience as the result of perceptual learning. The same processes involved in the perceptual learning of low-level properties are hypothesized to be involved in categorization (Tanaka and Taylor 1991; Goldstone and Steyvers 2001). Some of this categorization is of the sort that is involved in perceptual expertise.

1.4 Perceptual expertise

Perceptual expertise as it is studied in psychology is usually taken to refer to the ability of people to make rapid, automatic and reliable categorizations of objects. Most of us are proficient at making some categorizations – such as 'bird,' 'dog,' or 'car' – and so we possess some perceptual expertise ourselves. However, the perceptual experts that are typically the subjects of

psychological study are able to go a step beyond and make more fine-grained distinctions – they categorize birds as ‘northern thrush’ or ‘killdeer’, dogs as ‘Australian sheepdog’ or ‘lhasa apso,’ and cars as ‘Ford Model A’ or ‘1960s Jaguar E-type’.

Perceptual learning may be the means by which perceptual expertise is acquired. Perceptual expertise, like perceptual learning, also requires extensive exposure to exemplars of category members, and some of the same mechanisms, such as attentional weighting, are hypothesized to be involved in the development of perceptual expertise (Gauthier, Tarr, and Bub 2010; Kellman and Garrigan 2009; Kellman and Massey 2013; Y. K. Wong, Folstein, and Gauthier 2011). Here I call the sort of perceptual learning that results in perceptual expertise ‘perceptual learning for expertise.’ I focus primarily on this phenomenon in the work that follows.

Studies of perceptual expertise often involve comparing expert and novice processing and performance with a given class of objects. For example, novices are not able to distinguish between birds that trained ornithologists can easily differentiate (Tanaka, Curran, and Sheinberg 2005), and wine and beer experts are able to distinguish between wines and beers that novices cannot (Peron and Allen 1988; G. E. A. Solomon 1990). Through various experimental manipulations, several robust distinctions in the way such items are processed has been found. For example, experts process items of expertise via what has come to be labeled as configural or holistic processing, whereas novices do not (Gauthier and Nelson 2001; G. Rhodes et al. 1989; Gauthier et al. 2000). These processing differences, as I will argue in what follows, provide some insight into the high-level property debate.

1.5 Structure of this dissertation

I defend the perceptual expertise thesis: through perceptual learning experts come to perceive high-level properties imperceptible to novices.

Chapter two sets the stage by defending a certain kind of appeal to empirical data to settle the content debate. I term this general strategy ‘the common mechanism argument,’ and defend its use against objections, while also highlighting its weaknesses – the strategy must be used with caution.

My main defence of the perceptual expertise thesis is laid out in chapter three, where I employ the common mechanism argument, along with an appeal to the mechanisms of perceptual learning. In chapter four, I turn to aesthetic properties – which are high-level properties – and provide a novel view of how we can come to represent aesthetic properties in perceptual experience. This serves as an additional form of support for the perceptual expertise thesis.

The fifth and final argumentative chapter of my dissertation looks at the epistemic significance of perceptual learning. Here I argue that perceptual learning for expertise can provide us with a source of immediate justification for our perceptual beliefs.

1.5.1 Using empirical methods to adjudicate the content debate

I begin in chapter two by looking at previous attempts to appeal to empirical methods in order to resolve the content debate, and issues that they face. In the content debate, there has been an increasing trend towards appealing to empirical methods (Block 2014; Brogaard 2018;

Fish 2013; Block 2016; Bayne 2009; 2016; Briscoe 2015; Mandelbaum 2018; Stokes 2020).

Several of these appeals employ a similar strategy – a strategy I generalize in the ‘common mechanism argument.’ The strategy involves finding some relatively uncontroversial and well-studied effect or property (hereafter referred to only as ‘effect’) thought to be proprietary to the perceptual processing of low-level properties, and then arguing that it also applies to high-level properties, so such properties should be thought of as being represented in perception.

Such appeals, however, are usually made piecemeal, without proper attention to the various other empirical sources of conflict or support. Nothing is said about how to resolve conflicts between empirical studies that generate different, apparently contradictory results, and little is said to address potential confounds. In this chapter I provide some guidance as to how to navigate these issues in a way that strengthens usage and understanding of the common mechanism argument.

This chapter sets the stage for my application of the common mechanism argument to perceptual learning for expertise in chapter three.

1.5.2 What is learned in perceptual learning

In chapter three I provide my main defence of the perceptual expertise thesis. My argument proceeds in three parts. First, I appeal to the mechanics of perceptual learning to support the claim that we learn to represent high-level properties. Second, I employ what I term the ‘common mechanism argument’ in order to support the claim that such properties are represented in perceptual experience. Third, I draw on the phenomenon of categorical perception

to support the claim that such properties are represented in the phenomenal character of perceptual experience.

My argument for the claim that perceptual categories correspond to high-level properties appeals to the mechanics of perceptual learning. I argue that the perceptual learning processes of attentional weighting and topographical imprinting are what primarily allow for the construction of a (psychologically real) perceptual category. This perceptual category allows for the attribution of high-level properties to objects, via their categorization into what I call ‘perceptual kinds’ during perceptual processing.

I argue that this resulting perceptual category is best thought of as corresponding to a high-level property because it better accounts for the patterns encoded in the category than the post-perceptual judgment hypothesis, which claims that the attribution of high-level properties occurs post-perceptually. The systematic encoding of structural relations between low-level properties, as well as the evidential weighting of such properties is sufficient to diagnose in perception membership in a category that corresponds to a high-level property.

I then adapt the common mechanism argument to draw on studies of perceptual expertise. Many such studies involve comparing experts and novices in a given domain, or comparing a given population prior to and subsequent to training. Again the strategy involves finding an effect thought to be proprietary to perceptual processing. However, unlike the version reviewed in chapter two, this version of the strategy proceeds by positing that, prior to learning, our attributions of a given high-level property do not exhibit the relevant effect, and that afterwards they do. The best explanation for this is that they come to be represented in perception. I present the evidence for several such effects: speed, automaticity, encapsulation, and adaptation effects,

arguing that combined they support the claim that at least some high-level properties are represented in perception.

In the final section of this chapter, I provide an argument for the claim that such high-level properties are represented in perceptual experience. This might seem like overkill at first. One might think that from the arguments made in the previous sections I have already adequately defended this claim. However, whether this is the case depends on a further question as to how the contents of perception and the phenomenal character of perceptual experience (what it is like to have that experience) are related. If the contents of our perceptions can outstrip what our experiences are like, then it may be that while we represent high-level properties in perception, these properties have no impact on the phenomenal character of our perceptual experience.

Therefore, I provide an argument in favour of phenomenal richness that does not depend on adopting a particular theory of how perceptual representational content and the phenomenal character of perceptual experience are related. I argue that the phenomenon of categorical perception supports the claim that high-level properties can make a difference to the phenomenal character of perceptual experience.

1.5.3 Aesthetic perception

The last step of my defence of the perceptual expertise thesis is to turn to aesthetics in chapter four. Aesthetic properties such as beauty, ugliness, gracefulness and the like are high-level properties. While people have strong pre-theoretic intuitions that we perceive aesthetic properties, there may be some reason to reject a perceptual view. Aesthetic expertise seems to

require a great deal of training – we are not born wine connoisseurs or experts in female Renaissance painters. Instead, we acquire these abilities over time, through learning.

Training doesn't sit well with aesthetic for two reasons. First, it invites a disanalogy with our ability to perceive low-level properties such as size, colour, and shape. We do not seem to need training to perceive these properties, so if training is involved in the case of aesthetic properties then this may count against a perceptual view. Second, if we suppose that such training essentially involves learning about the artist, historical milieu, the production techniques used, and so on, then this suggests that we draw upon this knowledge somehow to make accurate aesthetic judgments, rather than simply perceive aesthetic properties.

Here I develop a positive view of how we can come to represent aesthetic properties in perceptual experience that accommodates the role of training, amongst other desiderata. This view builds on and develops themes from Kendall Walton's 1970 paper 'Categories of Art'. As such, I call the view 'Waltonian perceptualism.' Perceptual learning explains how we can come to represent new categories of art and objects of aesthetic appreciation in perception, and the relative fluency with which such objects are categorized explains how we can represent aesthetic properties in perceptual experience. An account of what it is to be an aesthetic expert provides us with the means of saying when those perceptual experiences are veridical, and the objects can truly be said to have those aesthetic properties.

1.6 The epistemology of perceptual learning

In chapter five I explore the epistemic significance of perceptual learning, and defend what I term the *wider scope* thesis: the scope of immediately justified perceptual beliefs can be

enlarged or broadened through learning. A belief is immediately justified when it does not depend on another belief for justification, but is nevertheless justified (Pryor 2005; Alston 1983; McGrath 2017). Mediatly justified beliefs are those that do depend on one or more other beliefs for justification. In the case of beliefs formed on the basis of perception (henceforth post-perceptual beliefs), the source of justification is not a further belief, but rather our perceptual experience.

However, the wider scope faces a challenge. If perceptual learning is a form of cognitive permeation, then it may not serve as a source of immediate justification. Cognitive permeation (also commonly referred to as cognitive penetration) is hypothesized to occur when the contents of perceptual experience are altered or permeated in some way by one's cognitive states, such as beliefs and desires (Pylyshyn 1999; Macpherson 2012; Stokes 2013; 2018b; Briscoe 2015; Raftopoulos 2009; Newen and Vetter 2017; Gross 2017; Siegel 2012; 2013).

To meet this challenge I defend the *pattern-driven* hypothesis: perceptual learning is driven by real patterns in the relevant class of training stimuli, and it is the gradual learning and detection of these patterns that drives the restructuring of the perceptual system, rather than declarative learning. I contrast it with the belief-driven hypothesis, which is an account of how declarative learning in the form of beliefs and (non-perceptual) concepts may drive perceptual learning in such a way as to constitute a case of cognitive permeation. Against the belief-driven hypothesis I argue that perceptual learning does not require cognitive guidance, and so does not count as a form of cognitive permeation. In cases where cognitive guidance is involved in perceptual learning, the sort of guidance does not suffice to count as cognitive permeation.

1.7 Significance of this work

In this dissertation I primarily focus on how perceptual learning for expertise can resolve the content debate in the philosophy of perception. Resolving the content debate is important in particular for the epistemology of perception. Perception is often thought to be a central source of foundational knowledge about the world, so the scope of the content of our perceptual experience will affect the scope of our foundational knowledge about the world.

However, the impact of this work is broader than the epistemology of perception. Here I explore three areas that go beyond philosophical debate. While I will not address these themes in depth in this dissertation, I include them so that readers may get a better sense of how this philosophical work can impact other fields.

1.7.1 Real-life expertise training

Our entire civilization is reliant on expertise. We depend on engineers, architects, and city planners to keep our cities functional. We depend on judges, lawyers and the police in order to keep our cities safe. We depend on doctors, nurses, medical technicians, chemists and pharmacists to keep us healthy. We depend on critics to direct us to the books, TV shows, plays and movies worth watching. While not all forms of expertise will involve a perceptual component, many will. The work in this dissertation furthers our understanding of the nature of perceptual expertise, and so provides guidance in designing more effective training programs in fields from radiology to the military to art criticism.

Understanding what is learned in terms of perceptual categories, and clarifying the role of cognition in guiding perceptual learning, both suggest that training in many disciplines ought to place a greater emphasis on exposure to a large (and unbiased) sample of exemplars of the category of interest, so that perceptual (and not just cognitive) expertise is achieved.

Furthermore, understanding what is learned in terms of perceptual categories serves to help clarify how perceptual learning for expertise can go wrong. That is, it provides a base structure from which we can begin to understand how perceptual learning can be biased, undergirding other biases and stereotypes, such as racial and gender-based biases. This research thus helps to make a positive contribution to emerging ‘debiasing’ training programs for professionals such as doctors, police officers and judges that have achieved mixed success (B. W. Smith and Slack 2015; Larrick 2004) (for further discussion of bias in perceptual learning for expertise see section 6.5.1).

1.7.2 Visual literacy

There is a growing call in academia to take seriously the idea of visual literacy across the disciplines: that just as we require training to learn to read written texts, so too do we require training to learn to ‘read’ images (Avgerinou and Ericson 1997). This call for training has become increasingly important given that the internet has vastly increased the number of

infographics, memes, photographs and other images that we are able to share in order to communicate.²

A better understanding of perceptual learning for expertise can help us to build effective visual literacy training programs in much the same way as it informs the development of real-life perceptual expertise. For example, understanding how aesthetic properties can come to be perceived may shed light on how images can be harnessed to manipulate public sentiment. My work on how we come to perceive aesthetic properties suggests affect can operate the perceptual domain; before cognition even begins, images may act upon our internal bodily states.

1.7.3 The perception-cognition divide

What do vision scientists study? Vision, of course. But what are the limits of what can be seen? Different scientists – including those studying other perceptual modalities as well – draw the divide between perception and cognition in different places. There is no universally accepted division. My work helps advance our understanding of where this division should be drawn (functionally, rather than anatomically).

In this dissertation I put forth an account of how categorization can take place within perception that does not rely on supposing that cognition is essential to the perceptual learning process. It provides a principled way of including some high-level properties as part of the contents of perceptual experience, rather than cognition, without itself relying on controversial

² In October 2019 I took part in an interdisciplinary research roundtable on visual literacy at the Peter Wall Institute at the University of British Columbia. I owe a debt of gratitude to the other participants for advancing my understanding of visual literacy and its importance. See <https://visual.pwias.ubc.ca> (accessed November 1st, 2019), for more information on this event and background on visual literacy.

background assumptions about where to draw the perception-cognition divide. While this work does not settle the boundaries of where to draw the divide, it does rule out divisions that exclude categorization or high-level properties from perceptual processing. It also demonstrates that it is not necessary to posit that cognitive permeation occurs in order to allow for such categorization and high-level property attribution to take place.

Chapter 2: Using Empirical Methods to Determine the Contents of Perception

Summary: Do we perceive only low-level properties such as shape, colour and motion, or do we also perceive high-level properties such as natural and artifactual kind properties, aesthetic, and moral properties? Recent debate on this question has largely been centered on the method of phenomenal contrast. In this chapter I evaluate the prospects of a newer strategy gaining momentum amongst philosophers. What I term the ‘common mechanism argument’ appeals to empirical evidence that a given perceptual effect common to low-level properties is also present with certain high-level properties, and that the best explanation for this is that high-level properties are also represented in perception. While the common mechanism argument avoids controversial appeals to intuition, it nevertheless confronts several challenges of its own. Here I evaluate the seriousness of such challenges, in each case arguing that they are surmountable, though pointing out the limitations of the strategy along the way.

2.1 Introduction

Whereas low-level properties such as shape, size, motion, and colour are generally agreed to form part of the contents of perception, it is controversial whether we perceive high-level properties such as artifactual or natural kind properties, aesthetic properties, and moral properties.³ This is known as the content debate. The position that we represent high-level properties in perceptual experience is sometimes known as the *rich content view*. The most discussed defence of this view is the method of phenomenal contrast (Siegel 2006b; 2010; Siewert 1998). Siegel's (2010) argument can be summarized as follows: grant the intuition that what it is like to see pine trees before and after one has gained the ability to recognize them is different. If there is a difference in the phenomenal character of the experience of the pine tree expert and novice, then there is a difference in the contents of their experience. This difference in content is best explained as a difference in terms of high-level properties – pine trees look different to experts because they now represent the property of being a pine tree in their perceptual experience.

However, those who deny that high-level properties can be represented in perceptual experience (the *sparse content view*) are unlikely to grant the “minimal intuition one has to have for the argument to get off the ground” (Siegel, 2010, p. 101). That is, they may deny that there is a difference in what it is like to see the relevant object upon developing the expert recognitional capacity (Dretske 2015; Lyons 2005). The change between novice and expert might be said instead to concern what sorts of inferences or associations the expert makes – while perceptual experience remains the same, other things change that explain the difference.

³ Here I set aside skeptical worries that there is no way of settling the issue (Lycan 2014) or that there is no fact of the matter as to whether we represent high-level properties in perception (Logue 2013).

Though Siegel has provided arguments against other such possibilities (2010, p. 101-113), there has also been much pushback (Butterfill 2009; Pautz 2009; Crutchfield 2012; Reiland 2014; Fürst 2017; Price 2009; O’Callaghan 2011; Koksvik 2015; Brogaard 2013). Such a persistent clash of intuitions may even be reason to mistrust the argument entirely, if the intuitions are being produced by one’s background theory (Fish 2013).

A more recent alternative proposal has been to appeal to empirical methods to settle the issue of whether we represent high-level properties in perceptual experience (Block 2014; Brogaard 2018; Fish 2013; Block 2016; Bayne 2009; 2016; Briscoe 2015; Mandelbaum 2018; Stokes 2020).⁴ Several of these appeals employ a common strategy – what I term here the ‘common mechanism argument.’⁵ While this turn to empirical methodology bypasses the intuition stalemate, it faces several new problems, which are surveyed here. I argue these problems are surmountable, however. When properly employed, the common mechanism argument provides multiple sources of evidence for a given high-level property, and is a plausible means of advancing the content debate.

2.2 The common mechanism argument

The common mechanism argument for determining whether high-level properties can be represented in perceptual experience involves finding some relatively uncontroversial and well-

⁴ Some of these authors focus on the contents of perception, rather than on the contents of perceptual experience. Here I take the stance that the deliverances of the perceptual system are what will give rise to the contents of perceptual experience (though there may be aspects of perception that remain unconscious and so do not contribute). While one might have a view that cognition or perceptual judgments also contribute to perceptual phenomenology, I think this confuses the matter at hand, and it is clearer to hold that any cognitive element may contribute to one’s overall (but not properly perceptual) phenomenology. So here my interest in the question of determining the contents of perceptual experience should be read as being importantly constrained by what the outputs of perceptual processes are.

⁵ The turn to empirical methodology goes beyond the common mechanism argument. For example (Bayne 2009; Block 2016).

studied effect or property (hereafter referred to only as ‘effect’) of the perceptual processing of low-level properties, and then arguing that it also applies to the processing of high-level properties. The idea behind this strategy is that the perceptual system – though most confine discussion to the visual system – is governed by a proprietary set of mechanisms or principles of operation, and it is these that give rise to the sorts of effects seen with low-level properties. If a certain class of high-level properties is found to exhibit the same effects, then the best explanation of this is that these properties are also represented in perceptual experience. That is, this constitutes evidence that they are governed by the same such proprietary perceptual mechanisms. The argument may be put as follows:

1. Low-level properties, which are represented in perceptual experience, exhibit effect x
2. High-level properties also exhibit effect x
3. The best explanation for (2) is that high-level properties are also represented in perceptual experience
4. Therefore, high-level properties are (probably) represented in perceptual experience

Reasoning along these lines serves as the rationale for the design of many of the scientific studies cited in this chapter.⁶ Philosophers have also increasingly appealed to the common mechanism argument. The effects most discussed by philosophers are adaptation, automaticity, encapsulation, attention and speed.

⁶ Not all embrace this rationale, however. For example, some instead assume that the high-level properties in question are represented in perception, and view their work as investigating whether high-level perceptual processing works by the same mechanisms as low-level perceptual processing. See for example (Webster and Macleod 2011).

2.2.1 Adaptation criterion

Adaptations are rapid and temporary adjustments in the sensitivity of our sensory systems that result in changes in perceptual awareness (Webster 2012). One well-known kind of adaptation concerns aftereffects, where prolonged exposure to one kind of stimulus causes a subsequent perceptual illusion of the perceptual ‘opposite’ of the stimulus, where this most commonly involves low-level properties. For example, adaptation to a direction of motion causes the well-known waterfall illusion, where the water in a static image of a waterfall is seen as moving in the opposite direction as the original direction of motion.⁷ Adaptation to a scene coloured one way causes a black and white version of the same scene to appear coloured in an opposite fashion (the original red areas are perceived as green, and so on).⁸ Several philosophers invoke the adaptation criterion as a means of advancing the content debate (Fish 2013; Block 2010; 2014; Brogaard 2018).

2.2.2 Automaticity criterion

Automaticity concerns whether or not an agent can prevent the processing of a class of stimuli. If the processing is automatic then, upon exposure to a stimulus, the system processes the information whether the agent wants to or not.⁹ Low-level properties are thought to be processed in such a manner: in the case of vision they are often taken to be the ‘visual primitives’

⁷ To see the illusion: <https://www.youtube.com/watch?v=oNhcpOIQCNs>

⁸ For low-level property adaptations, see for example (Durgin and Proffitt 1996; Suzuki and Cavanagh 1998; Clifford 2002).

⁹ See (Mandelbaum 2015) for a more fine-grained analysis of the automaticity of mental processes. Here I take it that perceptual processing is typically automatic in that given an input the processing must be initiated, and that it cannot be interrupted once initiated. However, lab techniques such as backward masking show that it is possible to interrupt perceptual processing, so this characterization is not wholly adequate.

that get computed in the early stages of processing. The only way the agent can interfere with this process is by cutting off the source of stimulus input, such as shutting her eyes. Stokes (2020) invokes the automaticity criterion to argue that objects of expertise – such as faces, cars, and birds – can come to be represented in perception. The expert, in contrast to the novice, cannot prevent herself from processing such objects in a holistic manner. Bayne (2009) also invokes automaticity to argue that natural and artifactual kind properties are represented in perceptual experience (though see Brogaard 2013).

2.2.3 Encapsulation criterion

Many theorists hold the view that at least some perceptual processing is informationally encapsulated: it is sensitive only to information contained within the perceptual system, along with perceptual input, and cannot be permeated by information we possess that is stored outside of this system, such as in cognition. While in the following discussion I equate informational encapsulation and cognitive impermeability, strictly speaking cognitive impermeability is only one form of informational encapsulation. For example, vision might be unencapsulated with respect to other sensory systems, such as is the case with multi-modal and cross modal effects, but nevertheless be encapsulated with respect to cognitive processes (so cognitively impermeable). The arguments I look at here focus on whether perceptual systems are encapsulated with respect to cognition.

Pylyshyn (1999) takes visual illusions as one line of evidence for the encapsulation of what he terms ‘early vision,’ as such illusions are generally insensitive to beliefs or evidence. For example, the Müller-Lyer illusion, where two lines are perceptually experienced as different in

length, despite the fact that we believe the two lines to be of equal length. Such encapsulation of vision may then be used as a criterion for supporting the claim that a property is represented in perceptual experience.¹⁰

Brogaard (2018) uses the encapsulation criterion as part of an argument that we hear meanings. She points to cases of certain sentences that people continue to find meaningful despite believing that they are nonsensical (Longworth 2008). In such cases, if the sentence meanings were processed by means of cognitive inferences, then they would be sensitive to such beliefs. Since they are not, then this is evidence against the claim that they are the products of cognitive inferences.

Bayne (2009) argues that object recognition is informationally encapsulated by appealing to research done with patients who have associative agnosia, where perception of low-level properties remains intact but the patient is unable to categorize the object. Providing such patients with beliefs about the category of the object via testimony does not generally affect recognition of the object.

2.2.4 Attention criterion

The attention criterion originates largely from evidence that low-level properties are processed in parallel, absent attention.¹¹ In this case, parallel processing is meant to be a feature of visual processing that applies uncontroversially to low-level properties; they are processed

¹⁰ This criterion is not open to those rich content theorists that hold that cognitive penetration occurs, whereby the contents of perceptual experience are said to be altered in a semantically coherent way by beliefs or other cognitive states (Siegel 2012). If a system is cognitively penetrable then it is not encapsulated.

¹¹ See for example (Braun and Julesz 1998; Treisman and Gelade 1980). But see (Carrasco 2011) for empirical work on the role of attention in early vision.

without attention. If evidence can be found for the parallel processing of high-level properties, and there is no alternative non-perceptual explanation for this, then they are genuinely perceptually represented.¹²

The evidence for parallel processing of low-level properties stems largely from visual search studies where subjects are tasked with searching for an object with a specific low-level property amongst an array of distractors – for example, a green circle amongst an array of red circles. Parallel processing is thought to occur when the time it takes to locate the target does not increase as the number of distractors increases – the time it takes to locate the green circle will remain roughly the same whether there are ten, twenty or one hundred red circles in the array.

The reason this is evidence for parallel processing is because, were attention involved, the expected outcome would be an increase in the time it takes to locate the target amidst distractors. Attentional resources are often thought to be limited, so the visual search must be conducted serially if attention is involved, with subjects making their way object by object through the array until they locate the target, increasing search time. When the time it takes to locate a target amidst an array of distractors does not depend on how many distractors there are, this is known as the ‘pop-out effect.’ The effect is well established for low-level properties. Simple colours, shapes and orientations, for example, pop out among distractors regardless of the number of distractors (Treisman and Gelade 1980; Wolfe, Cave, and Franzel 1989).¹³

¹² For example, proponent of the attentional criterion William Fish writes that “It would seem independently plausible, the thought goes, to suppose that, for processes of interpretation to take place, we would need to allocate additional cognitive resources to the task. This suggests a methodology for answering our question: if a property requires attention to be perceived, perhaps this is evidence that it should be located in the interpretative component of a visual experience; if it can be perceived preattentively, this is reason to think that it appears in the presentational component” (2013, 49). Note that Fish’s project is slightly different from the one put forth here. He is looking to determine whether a property belongs in the presentational or the interpretive component of visual experience. One might reasonably take the former to correspond to the phenomenal content of perceptual experience and the latter to correspond to any cognitive or conceptual content.

¹³ One caveat here is that search speed will be affected by how visually similar the distractor set is to the target. The more visually similar, the longer the search will take (Neisser, Novick, and Lazar 1963). This is why the pop out effect is not defined in terms of absolute reaction times, but in terms of the time being invariant with respect to the size of the distractor set.

Both (Block 2014) and (Brogaard 2018) appeal to the pop-out effect, and hence the attentional criterion, to adjudicate the content debate. Block discusses the example of human faces, which have been found to pop out amidst non-face distractors (Hershler and Hochstein 2005).¹⁴ Brogaard provides the example of the well-known ‘cocktail party effect,’ where one may hear one’s own name being uttered by someone in a crowded, noisy room despite many other conversations taking place (Cherry 1953; N. L. Wood and Cowan 1995).¹⁵

There are also other methods for testing for parallel processing. For example, subjects may be tasked with completing a central attentionally demanding task along with a peripheral task to detect a high-level property. It is hypothesized that successful completion of the peripheral task does not require attention, as the central task requires all one’s attentional resources (Prinz 2012, 118–22; but see Briscoe 2015). (Fish 2013) and (Bayne 2016) both refer to this sort of experiment, discussing the results of (Li et al. 2002). Another method – that (Bayne 2016; Fish 2013; Hershler and Hochstein 2005) all appeal to – involves measuring reaction and accuracy times for the categorization of one vs. two images containing high-level properties as targets (Rousselet, Fabre-Thorpe, and Thorpe 2002). Again, the idea is that speed and accuracy should remain constant between the two conditions if the processing is in parallel.

2.2.5 Speed criterion

¹⁴ However, several studies previous to this failed to find a pop-out effect for faces. See (Hershler and Hochstein 2005) for a discussion of why this might be the case.

¹⁵ Pop out effects have also been found for three dimensional features (Enns and Rensink 1991), depth from shading (Kleffner and Ramachandran 1992), and ‘membership in a contrasting race’ (Levin and Angelone 2001).

The speed criterion proceeds on the basis that speed matters because cognition is slower than perception. If recognition of high-level properties occurs only post-perceptually and involves cognition then it will be slower than perception; one must first perceive low-level shapes and colours, and then perform a cognitive inference to ascertain their identity as a given object such as a dog, or a table. Several philosophers invoke the speed criterion to argue for the representation of high-level properties in perceptual experience (Fodor 1983; Fish 2013; Bayne 2016; Brogaard 2018; Stokes 2020).

One measure of the speed of recognition of high-level properties is behavioural, where eye movement toward a target is recorded.¹⁶ For example, (Kirchner and Thorpe 2006) showed subjects two images of natural scenes simultaneously for 20 ms on a computer screen. While both were natural scenes, one contained an animal and one did not. Subjects were tasked with saccading (moving their eyes) to the image that contained an animal. On average, subjects could complete the task in 120 ms. Kirchner and Thorpe estimate that recognition occurs around 95-100 ms, with approximately 20 ms needed to prepare the saccade after recognition is achieved.

A way of getting more precise recognition timing information is to look at neural responses to particular features by way of single-cell recordings, where a microelectrode is placed, via invasive surgery, directly next to a brain cell and is then used to measure electrical activity in that cell in response to specific stimuli. However, this method is not thought to track conscious perceptual recognition, but rather the earliest detection or processing of a property by the perceptual system. Macaque monkeys performing discrimination tasks involving human

¹⁶ While there are studies that utilize other behavioral measures, such as ‘go/no go’ tasks where subjects must engage in a behavior such as pushing or releasing a button if target is detected (e.g. Fabre-Thorpe et al. (2001)), these studies face the additional hurdle of disentangling total response times from speed of recognition. Not only must subjects recognize the property, they must also select and execute the correct response. So there remains a question as to at what time recognition occurs, out of the total time to produce a response. For further experiments that look at natural object categorization speeds, see (Subramaniam, Biederman, and Madigan 2000, 511–35; Potter and Levy 1969, 10–15; Thorpe, Fize, and Marlot 1996, 520–22; Delorme, Richard, and Fabre-Thorpe 2000, 2187–2200; VanRullen and Thorpe 2001, 454–61; Thorpe et al. 2001, 869–76).

faces, monkeys, or food were found to exhibit neural responses in face-selective neurons within 20 to 40 ms (Nakamura, Mikami, and Kubota 1992). While human evidence is difficult to obtain because of the invasiveness of the technique, the procedure has been performed with epilepsy patients. In these patients category-specific neural firing has been found to occur within 100 ms for human faces, fruit, animals and chairs (H. Liu et al. 2009). To compare this data with low-level properties, human single-cell recordings with epileptic patients have found neural responses at 30-50 ms for certain simple visual stimuli such as flashing lights or moving patterns (Wilson et al. 1983).

A more common, non-invasive, method used on humans is electroencephalography (EEG), where electrodes placed along the scalp measure changes in the electrical activity generated by different brain regions. The averaged EEG responses are called event-related potentials (ERP). Human subjects exhibit differential ERP responses to novel vs. familiar faces within 50-70 ms (Seeck et al. 1997), to inverted vs. upright faces at 120 ms (Linkenkaer-Hansen et al. 1998), and for male vs. female faces between 50-65ms (Mouchetant-Rostaing, Giard, Bentin, et al. 2000; Mouchetant-Rostaing and Giard 2003; Mouchetant-Rostaing, Giard, Delpuech, et al. 2000). When subjects were asked to perform an animal detection task, differential ERP responses to scenes that contain an animal vs. those that do not were observed at 150 ms (Thorpe, Fize, and Marlot 1996).

In summary, there are a variety of effects one might draw upon when employing the common mechanism argument. However, this strategy faces several challenges. Here I will review what I take to be the three most pressing issues. In section 2.3 I respond to scepticism about the prospects of an appeal to empirical methods. In section 2.4 I respond to the criticism

that the strategy commits the converse fallacy. In section 2.5 I examine how well the empirical studies that the strategy employs can avoid low-level confounds. Finally, in section 2.6 I discuss three caveats to bear in mind when interpreting the results of the empirical studies: asymmetry of evidence, generalizing to a class, and conflicting results.

2.3 Scepticism about the prospects of empirical methodology

A preliminary objection to the project runs as follows: empirical considerations cannot help us adjudicate the content debate because science itself relies on theoretical background assumptions that require philosophical evaluation. Masrour (2011) expresses doubt of this sort. He points out that what constitutes the perceptual system is a substantive theoretical question. It cannot be answered via empirical methods alone – where to draw the divide between perception and cognition will depend on theoretical considerations such as what the function of the perceptual system is taken to be. And where one draws the boundaries of the perceptual system will in turn dictate whether or not high-level properties can be represented in perception.

However, the common mechanism argument does not require us to take a stand on where to draw the perception-cognition divide. We do not have to decide where perception ends and cognition begins prior to utilizing the strategy because we are relying instead on the uncontroversial assumption made by both parties to the debate that we perceive low-level properties. The argument proceeds from common ground – if the perceptual system processes anything, it is low-level properties.

The speed criterion is the place where this objection gains some purchase, but only for a particular set of cases. If high-level properties are found to be detected faster than or just as

rapidly as low-level properties, then we need not take a stand on when perceptual processing ends and cognition begins. This is because we are using the speed of the processing of low-level features as a benchmark. Likewise, if high-level properties were found to be detected only extremely slowly, then this would uncontroversially point to their being detected in cognition, rather than perception. This is because even rich content theorists will grant that very slow processing times implicate cognition.

The place where theoretical concerns are most relevant is with respect to the ‘contested middle’ – cases where high-level properties are processed fairly rapidly but nevertheless more slowly than low-level properties. Interpreting results in this grey area would seem to require taking a stand on when perceptual processing ends and cognitive processing begins, which in turn may depend on theoretical considerations and commitments. For example, if a certain processing stage involves drawing on long-term memory, then some might conclude that the process is cognitive, while others might conclude that some aspects of long-term memory are part of the visual system.¹⁷ So it would seem that the empirical results in the contested middle no longer straightforwardly favour one hypothesis over the other.

A further complication is that in order to determine exactly how quickly low-level properties are detected we must decide whether we are looking for evidence of the speed of the processing of individual low-level properties, or whether we are looking for evidence of the speed of the processing of conjunctions or ‘perceptual groupings’ of such low-level properties (e.g. ‘red square’ vs. ‘red’ and ‘square’). At first one might take the relevant comparison to be individual low-level properties – after all, we are comparing this to the speed of detection of

¹⁷ Pylyshyn endorses something like the former view. He writes “the early-vision system could encode any property whose identification does not require accessing general *memory*” (Pylyshyn, 2003, p. 136).

individual high-level properties, so this seems like an equal comparison. However, low-level theorists generally hold that we perceive groupings of low-level properties – low-level properties are bound together, and shapes are segregated into objects that appear to possess volume (Pylyshyn 1999, 343; 2003, 143; Raftopoulos 2009, 51). For this reason we might instead look to the speed of such processing as the relevant benchmark against which to measure the speed of high-level properties.

Which benchmark we choose makes a big difference. Human single-cell recordings with epileptic patients have found neural responses at 30-50 ms for certain simple visual stimuli such as flashing lights or moving patterns (Wilson et al. 1983). Conjunctions of low-level properties are processed more slowly, however (Bodelón, Fallah, and Reynolds 2007), and feature binding and object segregation have been hypothesized to occur during what is known as ‘local recurrent processing’, approximately 100-150 ms after stimulus onset (Lamme and Roelfsema 2000). If we use the latter speeds as the benchmark, then the contested middle will be processing that begins after approximately 150 ms. The exact cut-off point will be controversial, and likely depends on substantive theoretical commitments, but at least we may point to any case of processing in this ballpark as a candidate for perceptual processing – its speed does not definitively count against (or for) its inclusion in the contents of perceptual experience.

Finally, while the speed criterion on its own may be of little help when a property falls in the contested middle, it is noteworthy that the common mechanism argument has an important role to play in settling the theoretical debate, and may help us decide where the perception-cognition divide ought to be drawn. For example, suppose that a given high-level property is found to exhibit adaptation effects, automaticity, informational encapsulation, and parallel processing. If it is also detected at a reasonable speed (at some point within the contested middle)

then we might use these criteria as reasons to posit that cognition occurs only some time after this. More generally, rather than settling such theoretical issues independently of empirical evidence, we should favour a method of calibration. Theory will help guide empirical inquiry, but empirical results will also help guide theory.

2.4 The converse fallacy pitfall

The strength of the common mechanism argument depends in part on whether the relevant effect is truly proprietary to the perceptual system. If the cognitive system also produces the relevant effect or exhibits that characteristic, then this undermines the support for premise three of the common mechanism argument, that the best explanation for why high-level properties exhibit a given effect is that they are also represented in perceptual experience. The best explanation may be instead that they are processed by the cognitive system (see also Briscoe 2015 n. 7). This is an instance of the converse fallacy, as the reasoning may be spelled out as follows: if a property is perceptual, it exhibits characteristic x; high-level properties exhibit characteristic x; therefore, they are perceptual. But of course there may be another (non-perceptual) reason why high-level properties exhibit characteristic x, so the logic is fallacious.

Individually, several of the effects discussed above may fall prey to the pitfall. For example, while automaticity is a feature of perceptual processing, it is likely not an exclusive feature. Post-perceptual judgments – judgments formed largely on the basis of perceptual experience but which are not themselves part of perceptual experience – may also be formed automatically. Many of the non-perceptual beliefs we form are also candidates for automaticity:

if one entertains the beliefs that p and that if p then q , then it is quite difficult to prevent oneself from believing that q (though of course it can happen in some circumstances).¹⁸

Similarly, informational encapsulation is also probably not an exclusively perceptual feature. There may be informationally encapsulated modules that are not perceptual. For example, cheater detection has long been argued to be modular (Cosmides 1989; Cosmides et al. 2005; Van Lier, Revlin, and De Neys 2013). Öhman and Mineka (2001) posit a ‘fear module’ that is relatively encapsulated from ‘more advanced’ human cognition and is automatically activated by threatening stimuli. However, they do not posit that such a module is part of the perceptual system, instead conceiving it as its own module. At the limit are those who hold the view that all mental processes are modular (Carruthers 2006; Sperber 2001), where such modules are taken to be informationally encapsulated to varying degrees. If this is the case then the difference between perception and cognition will likely be one of degree of encapsulation.

Finally, the brain also engages in parallel processing beyond perception. While conscious thought is generally taken to proceed serially, one thought at a time, it is plausible that many cognitive operations take place in parallel, outside of the conscious arena. For example, there is some evidence that the cognitions needed for multi-tasking can take place in parallel, rather than serially, though this is still a subject of debate (Fischer and Plessow 2015).

While these effects are not proprietary to the perceptual system, they nevertheless provide some support for the perceptual hypothesis. First, this is because at least some cognitive processing is not automatic, not (or only very weakly) informationally encapsulated, and does take place serially. So demonstrating that high-level properties can sometimes be processed

¹⁸ See also Siegel (2011) for discussion of the automaticity of beliefs.

automatically, or be impervious to our beliefs, or be processed in parallel at least demonstrates that they are not the products of this type of cognition. For example, post-perceptual judgments are likely not informationally encapsulated, as there is ample evidence that they can be influenced by cognitive factors such as one's beliefs about the likely purpose of an experiment (Firestone and Scholl 2016). So the relevant contrast will be with the characteristics of post-perceptual judgments.

Moreover, more local, task-specific contrasts can still be informative. For example, while cognitions can occur very quickly when uncoupled from perception, the relevant task-specific context is the speed of cognition given a perceptual stimulus. In this context cognition will plausibly be slower than perception, because the stimulus must first be perceptually processed before it can be cognized. So the speed at which detection occurs is able to provide information that favours either a cognitive or perceptual interpretation.

A final issue is whether we have really identified the same effect in non-perceptual processing. That is, it may be that we have identified a cognitive effect that appears superficially similar to the perceptual effect, but that on closer inspection is different. If this can be shown with respect to any given effect, then it dissolves the worry of the converse fallacy. This may be the case for adaptation, which is often thought to be proprietary to the perceptual system (Block 2014).

Against this, (Helton 2016, 858) suggests that there might be conceptual adaptation, offering an example whereby browsing written descriptions of mansions causes subsequent descriptions of medium-sized houses to be judged as smaller. There are in fact many empirical results that uncover effects of this kind, known in social psychology as 'contrast effects.' Real-life cases of contrast effects are ubiquitous – a common marketing tactic is to put items on sale,

where comparison with the original price causes us to judge that the new price is cheap and reasonable, regardless of the actual value of the item.

However, contrast effects are arguably different from adaptation effects, despite the fact that the earliest accounts of contrast effects were grounded in perceptual adaptations, and their study in social psychology arose in part from the study of adaptation effects (Stapel and Suls 2007 ch. 1). Perceptual adaptation centrally involves a change in the way the adaptee looks to the subject, where this is the source of the change in judgment. While there have been many different theories of contrast effects, it appears that they largely (though perhaps not exclusively) involve a shift in decisional criteria that is not undergirded by a shift in perceptual experience (Stapel and Suls 2007 ch. 2). While on the one hand this seems like just what one would expect if one were looking for conceptual rather than perceptual adaptation, on the other hand it suggests a different locus for the effect. This is important because it suggests that the similarity is superficial – the processes may have similar effects on altering judgments but the cause is quite different. In the case of perceptual adaptation it is due to a difference in perceptual experience, and in the case of conceptual ‘adaptation’ it is due to a change in decisional criteria.

The point is underscored by considering that even the effects are not that similar. While perceptual adaptation is reliably produced through exposure to a perceptual adaptor, this is not the case with contrast effects. Here, exposure to a non-perceptual adaptor may produce a contrast effect, but also may produce an assimilation or priming effect. For example, in (Srull and Wyer 1979) subjects performed an initial task that involved unscrambling word sequences to create sentences. Some of these sentences described hostile behaviours. In a subsequent task, subjects read a vignette and were asked to rate the protagonist of the story on several dimensions, including hostility. Srull and Wyer found that ratings of hostility were higher for those subjects

who had unscrambled a higher number of sentences with hostile content, the opposite of what should happen in adaptation. The study was also replicated by using unconscious priming (Bargh and Pietromonaco 1982).

Such cases have led to explanations of assimilation and contrast that point to cognitive influences, such as how aware the subject is of the initial stimulus as a potentially biasing influence on their judgment, or the subject's beliefs about their relationship to the initial stimulus (Mussweiler 2001; Stapel and Suls 2007 ch. 1). For example, female subjects with low self-esteem who were asked to rate their own attractiveness after exposure to an attractive female face rated themselves lower if they were told they did not share the same birthday (a contrast effect), and higher if they were told that they did share the same birthday (an assimilation effect) (Brown et al. 1992). It is doubtful that such cognitive influences exert an effect in perceptual adaptation. Instead, the explanation of why perceptual adaptation occurs points to "changes in the response properties of neurons induced by the recent stimulus context" (Webster 2012, 1). Exposure to a stimulus reliably fatigues the neurons that detect that stimulus or its component features, causing an exaggerated response from neurons that detect the perceptual 'opposite' features.

Taken together, this evidence suggests that contrast effects should be distinguished from perceptual adaptation, despite their respective processes being similar in some ways. Nevertheless, there is a real worry that studies designed to reveal high-level adaptation effects may instead be measuring non-perceptual contrast effects on judgment, and better experimental methods are needed to rule out this possibility (Storrs 2015). So while adaptation avoids the converse fallacy, it is still important to make sure that studies are not misclassifying cases of contrast effects as cases of adaptation.

In summary, while many of the criteria reviewed here may suffer to some degree from the converse fallacy pitfall, the most charitable way of interpreting the empirical studies is that they are merely providing one more type of evidence for the hypothesis that such properties are represented in perceptual experience. It is not that automaticity or encapsulation serve on their own to settle the issue, but rather that, when combined with other criteria, they offer compelling evidence in favour of the perceptual hypothesis. The more effects that a given class of high-level properties exhibit, the stronger the reason to believe they are represented in perceptual experience.

2.5 Low-level confound pitfall

The success of the common mechanism argument will depend on whether there is a compelling case that the properties in question that exhibit the effect are best explained as high-level properties. Sometimes, the effects may be explained by low-level properties that are co-extensive with the relevant high-level properties. Such low-level properties then serve as a confound when interpreting the results. For example, at one time it was thought that whether something is a member of the (high-level) category ‘letter’ or ‘number’ was responsible for the pop-out effect of a target letter amongst an array of numbers, or number amongst an array of letters (Jonides and Gleitman 1972). But further scrutiny showed that in fact the (low-level) shape properties of letters are on average more similar to other letters than they are to the shape properties of numbers (and vice versa), and that this accounted for the pop-out effect (Krueger 1984). When the similarity of shape properties between numbers and letters was properly

controlled for, then the pop-out effect was eliminated, and so it was determined that the high-level property ‘number’ or ‘letter’ was not really responsible for the effect.

VanRullen (2006) argues that the face pop-out effect falls prey to the low-level confound pitfall, and that Hershler and Hochstein’s (2005) results are no more surprising than the finding that a red car will pop out amongst an array of distractors that are neither red nor cars. This does not suggest that the item pops out because it is a car, however, as the effect is wholly explicable in terms of its colour, a low-level property. Hershler and Hochstein’s findings of a pop out effect for faces can be explained in terms of low-level geometric properties that typically co-vary with faces rather than the high-level property of being a face. Faces typically have oval shapes, with the eyes, nose and mouth arranged in a T-shaped pattern, for example. The distractors that Hershler and Hochstein employ are visually dissimilar in that they do not all also possess these low-level properties.

Against this, Hershler and Hochstein (2006) argue that their original (2005) results count against the possibility of a low-level confound. In their first experiment they found a pop out effect for faces but not cars. In both instances, the distractors that they used were a heterogeneous mix of images depicting natural phenomena (animals, landscapes, faces in the car detection condition) and artificial phenomenal (other means of transport, cars in the face detection condition). While faces would differ from these other objects with respect to some of their low-level properties, every other image would also differ from each other with respect to some of their low-level properties as well. This makes the experiment unlike the case of detecting a number against a homogenous background of letters, or vice versa.

Moreover, the fact that cars did not pop out against the heterogeneous distractors despite having their own relatively unique set of low-level properties also counts against VanRullen’s

criticism. Their second (2005) experiment further underscores this point. Instead of heterogeneous distractors, Hershler and Hochstein used black and white line drawings of either faces or cars. The task was to either detect a face amongst a background of cars, or a car amongst a background of faces. Again, here they found an asymmetry, with a face popping out amongst a background of cars, but not vice versa. In this case, the low-level explanation is implausible. If it were true, then the low-level dissimilarity between cars and faces should equally facilitate a pop out effect for cars amongst a homogeneous background of faces. So it appears that at least in the case of the property of being a face, we have some reason to believe that it really is represented in perception on the attentional criterion.¹⁹

The low-level confound is also a worry for the speed criterion, where it relies on studies of object recognition that employ microelectrode implantation or use ERP responses. These techniques present special challenges insofar as speed of detection occurs prior to the time at which it is consciously perceived, so one cannot rely on behavioural measures to establish what is being detected. With respect to ERP responses, VanRullen and Thorpe explain the worry as follows: “if one was to ‘record’ electrical activity in a computer while it is processing an integer variable to determine whether it is a prime number, one would find that the least significant bit of the binary-encoded integer is on average more ‘active’ on prime than on nonprime numbers. However, it takes much more from the computer to decide if the variable is a prime number than to just check whether it is an odd or even integer. Before processing itself has even started, the preliminary encoding mechanisms can sometimes reflect, on average, high-level properties of the input variables” (2001b, p. 459).

¹⁹ See also (Webster and Macleod 2011, 1715) for discussion of the possibility of a low-level confound in the case of face adaptation, and (Block 2014) for a summary of experiments that count against the low-level confound hypothesis for face adaptation.

However, this worry can be addressed in most cases. One strategy (which can be generalized beyond the speed criterion) that experimenters employ to guard against potential low-level confounds is to try to vary the low-level properties while keeping the high-level properties constant. For instance, in microelectrode studies, to determine whether the neuron is responding to high-level features of a stimulus, neural responses are measured for the same object seen from different distances, as well as different viewpoints (Nakamura, Mikami, and Kubota 1992; H. Liu et al. 2009). This varies the low-level properties of the objects – for example, patches of colour will be smaller or larger in the first instance, or perhaps absent in the second depending on the angle and the object. If the neuron fires regardless of these sorts of low-level variations, then this is taken to be evidence that the detection is of the high-level property.

While this general strategy of varying low-level properties while keeping high-level properties constant may not be available in every experimental paradigm it is nevertheless applicable to a wide range of cases.²⁰ So it would seem that this pitfall is wholly avoidable so long as one relies on experiments that control for this sort of issue.

There is a further complication, however, that deserves greater philosophical discussion. While we may be able to empirically rule out low-level confounds of the sort described above, it may nevertheless be the case that whatever property we do represent in perception is not high-level. The sorts of considerations involved in determining what counts as a high-level property will not be empirical but theoretical, and so the common mechanism argument is not applicable in this instance. Several philosophers (Price 2009; Brogaard 2013) have argued that we cannot represent natural kind properties in perception based on ‘twin earth’ considerations. In brief,

²⁰ See (Webster 2011; 2012) for discussion of this technique for uncovering the locus of adaptation effects.

suppose that Oscar acquires the concept ‘tomato’ and his twin earth duplicate acquires the concept ‘twin tomato,’ where twin tomatoes look visually identical to our tomatoes, but are nevertheless not tomatoes because they differ molecularly – they are distinct natural kinds. Oscar and twin earth Oscar will plausibly both have the same phenomenological shift as a result of acquiring these concepts, on the assumption that such a shift does take place. However, this shift cannot be explained by positing that the acquisition of the concept causes both of them to represent the object as *a tomato* in perception. Twin earth Oscar has never been exposed to tomatoes, but instead twin earth tomatoes, and so it is not plausible to hold that he comes to represent the property of being a tomato (rather than the property of being a twin tomato) in perceptual experience. So we may conclude that we cannot represent natural kind properties in perceptual experience.

While assessing this argument is beyond the scope of this thesis, this is a case that highlights the limitations of the current strategy – it cannot directly help adjudicate this debate. Therefore, those who invoke the common mechanism argument for natural kind properties would do well to engage with this issue (as does Bayne 2016).

2.6 Interpreting the results

Here I survey three factors to keep in mind when interpreting the results of the common mechanism argument: asymmetry, generalizing to a class, and conflicting results.

2.6.1 Asymmetry

Sometimes the evidence is asymmetrical insofar as exhibiting an effect can count strongly in favour of a high-level property's being represented in perception, but failing to exhibit the effect does not count strongly against its being so represented, or vice versa. In general, automaticity and encapsulation are asymmetrical in favour of the sparse content theorist. That is, in cases where processing of a high-level property fails to exhibit these effects, this is strong evidence that it is not represented in perception; yet when processing of a property does exhibit one of these effects, this offers only weak evidence in favour of it's being represented in perception. The reason for this is that requiring conscious effort for processing and being influenced by beliefs are hallmark characteristics of cognition.²¹ For instance, lack of automatic processing would suggest that some sort of deliberative or inferential process is involved – some (ostensibly cognitive) effort or decision on the part of the subject is required in order to initiate or complete processing. So if processing of a property fails to exhibit automaticity, then this counts strongly against the property being represented in perception. However, if processing of the property does exhibit automaticity then this counts only weakly in favour because, as discussed above, some cognitive processing may be automatic.

Not all effects are asymmetric in favour of the sparse content theorist, however. Some run the other way. For instance, while the presence of parallel processing for a class of high-level properties seems to be a strong indicator that it is represented in perception (especially when combined with evidence of rapid processing), the absence of such processing may only be weak evidence that the class is cognitively inferred. This is because while parallel processing is thought to be a hallmark of early visual processing, visual search tasks that involve finding

²¹ Though some might argue that influence by beliefs is instead evidence for the cognitive penetration of perception.

conjunctions of low-level properties (e.g. ‘green square’) have been found to proceed serially (Treisman 1993).²² Sparse content theorists would be reluctant to say that green square objects are not part of the contents of perception – even they do not endorse a theory of the contents of perception quite that sparse.

Moreover there may be alternate explanations, compatible with their perceptual locus, for why high-level properties must be processed serially. For example, consider again the case of pop out during visual search for faces. While VanRullen (2006) argues that Hershler and Hochstein’s (2005) research falls prey to the low-level confound pitfall, he nevertheless holds that other research supports the claim that faces are represented in perception. In order to explain why there is no pop out effect for faces, despite their being represented in perception, he advances an alternate hypothesis: in instances such as visual search there will be “local competition for neuronal resources [...] within the large receptive fields underlying high-level processing of object categories” (VanRullen 2006, 3018).²³ This provides an apt illustration of how failure to find a pop-out effect for a given class of high-level properties is only relatively weak evidence for the sparse view. There may be an alternate explanation of why the effect is absent for a given class, one that does not preclude that the class is represented in perception.

The speed criterion is not clearly asymmetric for either side when the results fall in the contested middle. As previously noted, when the processing of high-level properties takes place as fast as or faster than low-level properties then this clearly favours the rich content view, and when the processing of high-level properties takes place extremely slowly then this favours the sparse content view. But when high-level properties are processed extremely quickly, though not

²² Though see (McElree and Carrasco 1999).

²³ See also (VanRullen, Reddy, and Koch 2004; VanRullen, Reddy, and Fei-Fei 2005).

as quickly as low-level properties, this result does not on its own provide clear evidence either way. This lack of asymmetry suggests that it is particularly important to combine measures of speed that fall into the contentious middle with the other criteria discussed here. These other results provide reasons to count the speed as further (albeit weak) evidence that the given properties are genuinely high-level.

2.6.2 Generalizing to a class

Another area for caution when interpreting the results is to what degree we are warranted in generalizing to a class based on empirical results. For example, Mandelbaum (2018) uses the speed criterion to support Fodor's (1983) suggestion – which itself is based on Rosch's work on categorization – that the outputs of perception are basic level categories.²⁴ Basic-level categories are defined as the most inclusive level of categorization for which category members share common perceptual attributes (Rosch et al. 1976). This definition is question-begging in the context of the debate over the contents of perception, and so I will understand the category more simply as one that exhibits a medium level of inclusivity. This class would include 'dog' and 'table,' for example. Superordinate categories are more abstract and will contain basic-level categories as members, though the members will generally not look very physically alike. For example, 'dog,' 'dolphin,' and 'pygmy shrew' will all be contained within the superordinate category 'animal.' The superordinate category of 'furniture' will contain 'table,' 'bookcase,' and 'lamp.' Subordinate categories are more fine-grained and less inclusive than basic-level

²⁴ Though he does not take his argument to bear directly on the contents of perceptual experience, because many of the studies he discusses are about the unconscious outputs of perception.

categories, and consist of different sub-types of the basic level. For example, the basic-level category ‘dog’ can be broken down into ‘Australian shepherd,’ ‘Boston terrier,’ and so on, where each constitutes a separate subordinate category. ‘Table’ can be further differentiated into ‘kitchen table,’ ‘dining table,’ ‘side table,’ and so on.

Rosch et al.’s (1976) experiments were meant to show in part that basic-level categories constitute the first type of categorization that individuals make when confronted with an image. She found that subjects were faster at categorizing images into their basic level, rather than subordinate or superordinate, categories. Mandelbaum (2018) discusses further studies meant to demonstrate that basic-level categorizations are performed too fast for cognition to make a contribution, proposing on this basis that only basic-level categories are represented in perception.

Mandelbaum’s proposal, however, faces several problems. Several studies have found that atypical members of a basic-level category (such as ‘penguin’ for the category ‘bird’) are categorized more slowly at the basic level than they are at the subordinate level (Murphy and Brownell 1985; Jolicoeur, Gluck, and Kosslyn 1984), and that contextual cues can help subjects make subordinate level categorizations just as quickly as basic level ones (Murphy and Wisniewski 1989). Another study demonstrates that superordinate categories (e.g. animal) can also be identified more rapidly than basic-level categories (e.g. dog) in natural scenes (Macé et al. 2009).²⁵

Reflecting on research on perceptual expertise suggests that category-wide generalizations of the subordinate/superordinate/basic-level variety are too general to be useful.

²⁵ See also (Greene and Oliva 2009; Oliva and Torralba 2001) for arguments that superordinate categories are useful for facilitating basic-level categorization.

What matters is the individual's expertise with the category (Gauthier, Tarr, and Bub 2010). Though there will be many shared areas of proficiency in a culture (perhaps most commonly with items that fall into the basic-level category), there will also be individual variation.

While the subjects tested in Rosch et al.'s and subsequent experiments discussed above were not experts with subordinate categories, Tanaka and Taylor (1991) looked at such cases by studying bird and dog experts. They found that for these experts subordinate level categorizations in their respective areas of expertise (e.g. robin, German shepherd) were just as fast as basic-level categorizations, and in fact faster when it came to denying membership in a class (indicating that an image is not a robin vs. that it is not a bird). Tanaka and Taylor take their results to be consistent with Murphy and Brownell's (1985) 'differentiation hypothesis,' which holds that the speed of categorization will depend on how differentiated the category is for the subject. Such differentiation will involve learning the specific and distinctive features of the category. It seems then that we ought to proceed cautiously and extend our conclusions to the narrowest subclass needed to explain the results.

2.6.3 Conflicting results

A final issue of central importance to the common mechanism argument is the significance of conflicting results. For example, Hershler and Hochstein (2005) found that animal faces are not processed in parallel, but rather serially. On the attentional criterion, they would therefore be denied perceptual representation status. However, on the speed criterion animal faces would be accorded such status. It has been found that animal faces are detected just

as rapidly as human faces in the RSVP (rapid serial visual search) paradigm (Rousselet, Macé, and Fabre-Thorpe 2003).

What should we make of these conflicts? Should they make us sceptical of the strategy itself? Hershler and Hochstein briefly consider this conflicting data, writing that it “seems to indicate a different underlying process for parallel human face search and rapid serial picture identification” (2005, p 1718). It is hard to know what to make of this suggestion. On the most radical interpretation, the suggestion is that a given high-level property is represented in perception when the perceptual system is engaging in certain tasks, but not in others, despite detection of the property being relevant for both tasks. This would signify that the contents of perception are much more context and task-dependent than previously thought. However, as we have seen in the case of perceptual expertise, perceptual contents may change with the learning of subordinate categories, so perhaps this idea is not so radical.

It also may be the case that conflicting results can be found for low-level properties within a criterion. For example, subjects can detect a target colour without attention in the visual search task described in section 2.2.4, where subjects must identify, say, a red circle amongst an array of green circles (Treisman and Gelade 1980; Wolfe, Cave, and Franzel 1989). In such cases the colour ‘pops out’ at the subject quite rapidly regardless of the size of the array, and so is taken to signal that colour processing can occur in parallel. However, this verdict appears to conflict with the results of another experiment, where subjects perform an attentionally demanding central task while also performing a peripheral colour detection task (Li et al. 2002). The peripheral task is to detect a circle with a red half on the right and a green half on the left. The distractor is a circle in which the colours were reversed. In this case, subjects could not perform this peripheral task above chance in the absence of attention.

Perhaps there is enough of a difference between these two tasks so that we can explain away the apparent conflict – for example, maybe a peripheral task that involved the detection of a simple green or red circle would not require attention. However, even if we suppose the conflict is genuine, then instead of viewing this as grounds for scepticism about the common mechanism argument we should instead invoke asymmetry. A positive result (detection in the absence of attention) is strong evidence that the property is represented in perception, but a negative result, at least in this case, should not be counted as evidence against the perceptual view, given the bulk of evidence that we represent low-level properties in perception. In general, asymmetry may be relevant to settling conflicting results in this way.

2.7 Conclusion

Here I have provided an overview of a promising means of advancing the debate on the contents of perception – the common mechanism argument. It involves finding a property or effect of the perceptual processing of low-level properties and testing to see whether high-level properties also exhibit the property or effect. Two advantages of this strategy are that it does not rely on philosophical intuitions about phenomenology, and that it begins from common ground between sparse and rich content theorists: both sides agree that the processing of low-level properties is perceptual. I have argued that this second advantage is what allows the strategy to avoid the criticism that empirical methodology cannot settle the content debate because such methodology itself depends on theoretical commitments.

While some of the criteria do fall prey to the converse fallacy in that the perceptual effect in question is not proprietary to the perceptual system, the strategy is nevertheless useful for two

reasons. First because its strength lies in demonstrating that the processing of high-level properties exhibits several of the effects. Second, because some effects are proprietary to the perceptual system when it comes to specific tasks, and so the results of these experiments are informative.

I have also argued that while there are sound experimental techniques for overcoming one version of the low-level confound issue, the common mechanism argument does not have the resources to address a more substantive philosophical version of the challenge to natural kind properties involving twin earth considerations. This highlights a limitation of the strategy, and suggests that more philosophical work ought to be done to address this challenge.

Finally, I have highlighted three issues that arise when interpreting the results of studies employed by the common mechanism argument. First, there may be an evidential asymmetry for a given effect, where if the processing of a high-level property exhibits a given effect this counts strongly in favour of its being represented in perception, but if the processing does not exhibit the effect then this counts only weakly against it being represented in perception (or vice versa). Second, we ought to generalize to the narrowest subclass of properties needed in order to explain the empirical results, in part because studies of perceptual expertise suggest that the contents of perception will alter with experience. Third, in the face of conflicting results we may nevertheless favour the positive findings (that a high-level property is represented in perception) because there may be alternate task-specific explanations of why processing of the high-level property does not exhibit the effect, compatible with its being represented in perceptual experience. Asymmetry may also be useful in interpreting conflicting results.

In summary, the common mechanism argument represents a powerful tool for adjudicating the content debate when multiple criteria are invoked with respect to a given high-

level property, though care must be taken to understand the limitations of the strategy, and to interpret the results.

Chapter 3: What is Learned in Perceptual Learning

Summary: Perceptual learning is an enduring change in the perceptual system – and our resulting perceptions – due to practice or repeated exposure to a perceptual stimulus. But what exactly is learned in perceptual learning? Does our perceptual experience become richer only in virtue of representing more low-level properties such as colours, shapes, and motion, or does it also come to include the representation of high-level properties such as dogs (natural kinds) and toasters (artifactual kinds)? Here I argue that we can come to represent high-level properties in perceptual experience, where this has an effect on the experience's phenomenal character. My argument proceeds in three parts. First, I appeal to the mechanics of perceptual learning to support the claim that we learn to represent high-level properties. Second, I employ what I term the 'common mechanism argument' in order to support the claim that such properties are represented in perceptual experience. Third, I draw on the phenomenon of categorical perception to support the claim that such properties are represented in the phenomenal character of perceptual experience.

3.1 Introduction

Perceptual learning is an enduring change in the perceptual system – and our resulting perceptions – due to practice or repeated exposure to a perceptual stimulus. While it is part of infant development, adults are also capable of undergoing perceptual learning (Gilbert, Sigman, and Crist 2001), and it is hypothesized to be a component of many forms of specialized expertise (Kellman and Garrigan 2009). But what exactly is learned, or what can be learned, in perceptual learning? Eleanor and James Gibson posed this as a central question for the study of the phenomenon (J. Gibson and Gibson 1955b). Another way of asking this question is in terms of whether perceptual learning changes the contents of perceptual experience, and if so, how.

To say that perceptual experiences have contents is to say that such experiences can be assessed for accuracy – they convey to the perceiver that the world is a certain way, where what is conveyed can either be accurate or inaccurate. One way in which these contents can be assessed for accuracy is in terms of the properties they ascribe – truly or falsely – to objects. The content debate concerns what sorts of properties can be part of perceptual experience. Those who hold the position that we perceptually experience only shapes, colours, motion, relative position, and the like are sparse content theorists. They hold that we perceptually experience only low-level properties, such as the ones listed above, along with analogues in the other sensory modalities (Tye 1995; 2000; Dretske 1981; Clark 2000; Raftopoulos 2009; Connolly 2014; 2019). Low-level properties are usually defined stipulatively – they are the properties that everyone agrees are represented in perceptual experience. High-level properties are all those properties that are not low-level. They include natural and artifactual kind properties, aesthetic and moral properties, and properties such as causation. Those who hold the position that we

perceptually experience at least some high-level properties are rich content theorists (Siegel 2010; 2006b; Fish 2013; Werner 2016; Siewert 1998; Brogaard 2018; Stokes 2018a).

The question of what is learned in perceptual learning can thus be rephrased in the terms of the content debate: in perceptual learning can we learn only to represent more (or more fine-grained) low-level properties such as colours, shapes, flavours, and pitches, or can we also learn to represent high-level properties?

A further question can be posed – also informed by the content debate – as to whether what is learned makes a difference to the phenomenal character of perceptual experience. On certain theories of the relationship between representational content and the phenomenal character of perceptual experience, the representational content of perceptual experience will always exhaustively determine its phenomenal character (Tye 2002). However, one might hold the view that the representational content of perceptual experience does not always impact phenomenal character – the contents of perceptual experience can outstrip what our experiences are like. If this is the case, then it may be that while we represent high-level properties in perception, these properties have no impact on the phenomenal character of our perceptual experience (Prinz 2013; Chudnoff 2018).

Here I will answer both questions as follows: perceptual learning results in the representation of high level properties in perceptual experience, where this is understood to impact the phenomenal character perceptual experience.

3.1.1 Why does the debate matter?

The answer to these questions is important in part because perceptual learning provides a new way of making progress in the content debate. It has become common for rich content theorists to appeal to differences between experts and novices in order to argue that we can come to represent high-level properties in perceptual experience. However, rather than turning to empirical studies of such experts, appeals to expertise usually invoke intuitions that there is a change in the perceptual experience upon acquiring expertise. For example, Susanna Siegel argues, in what has come to be called the phenomenal contrast argument, that there is a difference in what it is like to see a pine tree or some such object of expertise before and after learning takes place, and that this difference is best explained by the representation of the high-level property of being a pine tree in perceptual experience (Siegel 2010; see also Siewert 1998; Stokes 2018a; Werner 2016).

However, those who deny that high-level properties can be represented in perceptual experience are unlikely to grant the “minimal intuition one has to have for the argument to get off the ground” (Siegel 2010, p. 101). That is, they may deny that there is a difference in what it is like to perceive the relevant object upon developing the capacity to identify the relevant property (Dretske 2015; Lyons 2005). Such a clash of intuitions may be reason to mistrust the argument entirely, if the intuitions are being produced by one’s background theory of perception (Fish 2013). In the face of intractable disagreement, some philosophers have argued recently that the contents of perceptual experience are vague, or that there is no fact of the matter (Logue 2013; Lycan 2014).

Such drastic conclusions are premature. Perceptual learning can help us move past the intuition stalemate by independently motivating the claim that there is a difference in perceptual experience between an agent prior to developing a recognitional capacity (the novice) and after

(the expert), and that this difference is best explained in terms of the representation of high-level properties. In this chapter I defend the perceptual expertise thesis: through perceptual learning experts come to represent high-level properties in perceptual experience imperceptible to novices. My argument proceeds in three parts. First I appeal to the mechanics of perceptual learning to support the claim that we learn to represent *high-level properties*. Second I appeal to what I term the ‘common mechanism argument’ in order to support the claim that such properties are represented in *perceptual experience*. And third I appeal to a phenomenon called ‘categorical perception’ in order to support the claim that such properties are represented in the *phenomenal character* of perceptual experience.

3.1.2 What is perceptual learning?

Perceptual learning can be characterized as an enduring change in the perceptual system or perceptual processing due to practice or repeated exposure to a perceptual stimulus (Goldstone 1998; E. J. Gibson 1963). Here, the perceptual system can be understood as whatever cognitive resources are causally responsible for producing perceptual experience. This change in the perceptual system usually allows subjects to make new perceptual discriminations, and this new ability is hypothesized to be undergirded by a change in perceptual experience – perceptual learning alters what it is like to perceive the world.

In some studies of perceptual learning, what is learned is likely wholly explicable in terms of low-level properties. For example, a common measure of perceptual learning called the Vernier acuity task involves judging whether the lower portion of a line is to the left, right, or

aligned with, the upper portion of the line (Westheimer and Hauske 1975; Mckee and Westheimer 1978). The ability to make such discriminations improves with practice. Likewise for other common perceptual learning tasks: training improves subjects' ability to discriminate the direction of motion of dots on a screen (Ball and Sekuler 1987; 1982; Matthews et al. 1999), spatial frequencies (Fiorentini and Berardi 1980), discriminate targets from background based on orientation (Karni and Sagi 1991), see form from motion (Vidyasagar and Stuart 1993), spatiotemporal interpolation (De Luca and Fahle 1999), and stereoscopic depth perception (Ramachandran and Braddick 1973; Westheimer and McKee 1978; O'Toole and Kersten 1992; Sowden et al. 1996). In all such cases, subjects are able to make discriminations that they were not able to prior to training. However, this improvement can be explained by positing that one's perceptual experience has become richer only with respect to low-level properties: we become better perceivers of shapes and their relations to each other, to the true direction of motion, and so on.

If this were all there is to perceptual learning, then it might seem that the answer to the question of what is learned in perceptual learning ought to be answered in favour of the sparse content theorist. However, this is not the end of the story. Our ability to categorize objects – where some of these categories correspond to high-level properties – is in some cases the result of perceptual learning. Going forward, I will call these categories acquired via perceptual learning 'perceptual categories.'²⁶ I understand perceptual categories to be psychologically real entities that are used in perceptual processing to categorize objects. In some cases, these perceptual categories correspond to high-level properties: once the perceptual system has

²⁶ While there may also be perceptual categories that are innate rather than learned, here I focus on those perceptual categories acquired via perceptual learning. It is also worth noting that I avoid calling these entities concepts for reasons that I will discuss in chapter five – at this stage of investigation we should not presuppose that perceptual categories are concepts.

determined that an object belongs in a given perceptual category, this allows for the attribution the relevant high-level property to that object, and the high-level property is represented in perceptual experience. I understand the act of perceptual categorization as that of categorizing objects in perceptual processing via the application of perceptual categories.

Following the work of Rosch et al. (1976), categories can be divided into three main nested types that are more or less inclusive. While these distinctions themselves do not presuppose that any of these categories are perceptual, there is evidence, reviewed in chapter two and section 3.3 of this chapter, that properties of each type can come to be represented in perceptual experience through perceptual learning. Basic-level categories are those that correspond to a medium level of inclusivity. These are categories such as ‘dog’ and ‘table.’ Superordinate categories are the most inclusive and tend to be quite abstract, in the sense that there is often no characteristic look to all category members (though this need not be the case). This will include the category ‘furniture,’ ‘animals,’ and so on. Subordinate categories are less inclusive. These include categories such as ‘German shepherd,’ ‘Boston terrier,’ ‘side table,’ ‘kitchen table,’ and so on. While the category of dog includes all Boston terriers, the category of Boston terrier does not include all dogs.

The ability to make subordinate categorizations is sometimes studied under the umbrella term ‘perceptual expertise’ (Gauthier, Tarr, and Bub 2010).²⁷ Perceptual expertise refers to a certain class of acquired recognitional abilities that psychologists have thought of as perceptual. The most commonly studied groups of real-world perceptual experts are radiologists, birdwatchers and ornithologists, chess players, dog fanciers, and car connoisseurs. Such people

²⁷ While perceptual learning and perceptual expertise have become relatively independent fields of research concerned with the learning of low and high-level properties, respectively, see (Y. K. Wong, Folstein, and Gauthier 2011) for empirical work bridging the two fields, and discussion in (Kellman and Garrigan 2009; Kellman and Massey 2013; Goldstone and Barsalou 1998; Goldstone, Landy, and Brunel 2011).

are able to make very rapid and reliable high-level categorizations of the objects of their expertise, unlike novices. Perceptual expertise can also be cultivated in the lab, either with real-life objects such as cars or with novel lab-created classes of objects such as ‘greebles’ and ‘ziggerins.’ And finally, perceptual expertise studies include some objects of expertise that (almost) all of us are experts in, such as faces. The study of perceptual expertise thus involves the study of our perceptual categorization abilities.

While psychologists call such expertise ‘perceptual,’ we might wonder whether it truly is, or whether it is instead a learned cognitive ability to make rapid inferences or judgments. I will return to this question in section 3.3. However, first I will argue that perceptual categories correspond to high-level properties in some cases.

3.2 Perceptual categories are high-level

Perceptual expertise involves extensive training with many exemplars of category members. In the lab, training people to become proficient at categorizing novel classes of objects can be achieved in approximately nine hours, through hundreds of categorization trials (Gauthier and Tarr 1997; Gauthier et al. 1998). Real-life experts at identifying subordinate categories of birds, dogs, tumours, wines, and so on spend years engaging with exemplars of the categories they are interested in. The perceptual category is therefore plausibly formed with the purpose of detecting members of that category. This alone should perhaps push us to the conclusion that the perceptual categories are high-level: such categories represent what they are set up to detect (see (Prinz 2006) for an argument along these lines). However, this fails to rule out an alternate low-

level explanation of the acquired abilities of experts. Several philosophers have put forth just such explanations.

3.2.1 The post-perceptual judgment hypothesis

The post-perceptual judgment hypothesis (hereafter abbreviated to ‘judgment hypothesis’) consists of two central claims: (1) the attribution of high-level properties occurs post-perceptually, involving some sort of judgment, concept-application or inferential process; and (2) any heightened perceptual abilities of experts can be explained without appealing to high-level properties. I will address the first claim further in section 3.3, and here focus on the second claim.

Those who advance some version of the judgment hypothesis often explain the heightened abilities of experts by positing that they acquire distinct learned patterns of attention (Pylyshyn 1999; Connolly 2014, 2019; Price 2009). As Zenon Pylyshyn writes, “An expert’s perceptual skill frequently differs from a beginner’s in that the expert has learned where the critical distinguishing information is located within the stimulus pattern. In that case the expert can direct focal attention to the critical locations, allowing the independent visual process to do the rest” (Pylyshyn 1999, 359).

Adding to this, Tyler Burge (2014) proposes that perceptual expertise allows us to form new perceptual groupings of low-level properties via learned attentional patterns that then form the basis for post-perceptual judgments of high-level properties. While Burge speaks in terms of ‘attributives’ rather than properties – where these are understood as syntactic vehicles that

attribute properties (Burge 2010, 32–33) – his proposal can be adapted to be understood directly in terms of properties. Burge writes:

in the acquisition of expertise, conceptual higher-level attributives can cause attention to and new grouping of specific low-level attributes. For example, expertise in bird species or piano types can produce a capacity to attend to, group, distinguish, and remember certain geometrical patterns, in ways that a non-expert would not. It is known that practice and expertise can affect organization of perceptual capacities for grouping shapes. In such cases expertise-guided attention to specific low-level attributes can help form generic low-level attributives that constitute a refined perceptual applicational basis for non-perceptual higher-level attributives. No new, perceptual higher-level attributives need be acquired.

(2014, p.10)

On this proposal for filling out the judgment hypothesis, our attention is habitually tuned to seek out certain low-level properties, and this eventually results in the formation of a perceptual category understood as a psychological entity that – in contrast to the view I defend in this chapter – allows us to detect a mere assemblage of low-level properties. Detection of this mere assemblage of low-level properties then forms the basis of a post-perceptual judgment that a certain high-level property is instantiated. The judgment, while made mostly on the basis of perceptual information, is nevertheless cognitively executed. The judgment hypothesis can also explain why our perceptual categories are so systematically correlated with the ascription of high-level properties. Expertise is a matter of gathering perceptual evidence via the perceptual

category that is then assessed in post-perceptual judgment, where the high-level property is ascribed.

How do we then adjudicate between these two hypotheses? On the perceptual expertise hypothesis, perceptual categories correspond to high-level properties. On the judgment hypothesis, they correspond to mere assemblages of low-level properties. In order to proceed, it is important to get clearer on what counts as a high-level property.

3.2.2 What counts as a high-level property

It has been customary in the debate over the contents of perception to simply stipulate that properties such as colour, shape, and motion are low-level properties, and then define high-level properties as those that are not included in the list of low-level properties. The list is stipulative in part because it is quite difficult to come up with a principled way of making the classification. Instead, the impetus behind the groupings is that low-level properties are those that most everyone agrees are part of the contents of experience, and high-level properties are the controversial ones.

Understood in this way, so as long our perceptual categories do not correspond to mere assemblages of low-level properties, then they should be considered as corresponding to high-level properties. It is thus a further question whether our perceptual categories correspond in some cases to natural kind properties. Whether or not this is the case depends in part on what we take natural kinds to be, and so is a question that goes beyond this thesis (Brogaard 2013; Bayne 2016). In order to remain agnostic on the issue here, I will call the properties that correspond to some perceptual categories ‘perceptual kinds.’

Perceptual kinds may or may not count as natural kinds, but they are high-level properties. Moreover, many of the perceptual kind properties tracked by our perceptual categories are importantly linked to natural kinds – in some scenarios they will track such kinds almost perfectly, and if they track them less than perfectly there is still significant overlap. While we may get it wrong sometimes, we are mostly able to perceptually categorize dogs as dogs (understood here as a perceptual kind). If perceptual kinds do not qualify as natural kinds themselves, then they certainly play an important role in our discovery of natural kinds.

(Consider: without some initially salient perceptual similarities would we ever have arrived at the scientific taxonomies we have now? Perceptual similarity serves as a departure point in our scientific theorizing. While there are probably infinite ways of grouping objects in terms of perceptual similarities, there are nevertheless certain ways of grouping objects that are more useful for our purposes. So the actual perceptual kinds that we learn to represent in perceptual experience will be a fraction of the perceptual kinds out there. If we posit that grouping objects in terms of natural kinds will be useful for our purposes sometimes, and that this can be accomplished in some cases by relying on perceptual similarities, then our perceptual kinds will track natural kinds.)

The relevant question going forward will then be whether perceptual kinds are high-level properties.

3.2.3 Argument from the mechanics of perceptual learning

The alternative to viewing perceptual kinds as high-level properties is that they are mere assemblages of low-level properties, along the lines proposed by Burge. An alternative to

positing perceptual kinds at all is that experts are simply better at attending to the locations likely to harbour relevant information for categorization, along the lines proposed by Pylyshyn. Here I appeal to the mechanics of perceptual learning to argue against these views.

The argument begins with an understanding of how, beyond exposure to many exemplars of category members, the formation of perceptual categories via perceptual learning occurs.²⁸ This is the subject of much empirical study, and several different models of object categorization have been proposed (Palmeri and Cottrell 2010). However, many such models make use of the idea of selective attention, or attentional weighting.

In attentional weighting, an object's various properties or features are 'weighted' more or less heavily, depending on their usefulness for performing tasks, especially those tasks that produce a reward (Goldstone 1998; Goldstone and Steyvers 2001). If a stimulus feature has been useful in the past for predicting rewards, then that property will tend to be weighted more heavily.

While such weighting has been studied with respect to low-level perceptual learning (Doshier and Lu 1998; J. Liu, Doshier, and Lu 2015), it is also hypothesized to play a role in perceptual categorization (Tanaka and Taylor 1991; Goldstone and Steyvers 2001). While some tasks are transitory, categorization is a stable task – we categorize objects by default, across contexts, without external prompts. This reflects the fact that categorization is almost always important, regardless of context – it allows us to access information about potential threats and rewards in our immediate environment, amongst other potentially relevant information.

Attentional weighting corresponds to a learned pattern of feature selection; it guides us to

²⁸ Here I put aside unitization and differentiation, two other components of perceptual learning (Goldstone 1998). While they are also important to the learning of perceptual categories, they do not enter into the argument I provide.

highlight in perceptual processing – and ultimately in perceptual experience – those features that are useful for the task. In categorization, those features that are useful will be those that are highly diagnostic of category membership. A robin’s brightly coloured red breast is highly diagnostic of its being a robin, for example.

Here we can contrast the novice at performing some task with the expert. While the novice is uncertain as to which aspects of a visual stimulus or array are or will be relevant to the task, the expert knows where to look, and what to look at. Attention is tuned relative to the task demands. For example, radiologists, who are experts at detecting tumours via x-rays (Drew et al. 2013; van der Gijp et al. 2017), and artists, who are experts at assessing the composition of paintings along with other aesthetic features (Vogt and Magnussen 2007), both exhibit attentional patterns that are systematically different from those of novices.

Explained this way, it might seem that what the expert learns in perception is a mere attentional skill, and so attentional weighting supports the simplest version of the judgment hypothesis: we learn to attend more quickly to the areas most relevant for categorization, and this information is fed to the cognitive system in order to determine whether something is a category member.

However, this picture is mistaken on several fronts. Attentional weighting should not be thought of as a mere orienting of attentional resources. First, the weighting is not with respect to location, but to the property itself.²⁹ Attention to properties, as opposed to spatial locations, allows it to be the case that the novice and the expert can both be looking at the same spatial

²⁹ For a review of feature-based attention see (Carrasco 2011).

location but attending to different properties. In this sense it is not true that the expert simply knows where to look – her attention is also guided to the relevant properties.

Second, attentional weighting is not a mere orienting of attentional resources because it also provides a measure of significance or importance of the feature for the organism. The information contained in the learned attentional weights is not simply ‘look at this property’ but is rather best understood as something akin to ‘this property is important to x degree in confirming or rejecting this exemplar as a member of a given category y. Given that once perceptual learning has occurred, attentional weighting becomes relatively stable for categorization (Richler, Wong, and Gauthier 2011), we should consider these weights as directing attention without the need for cognitive guidance, in accordance with how diagnostic the property is for categorization. The significance of the attended property is already encoded in the attentional weight itself. It therefore goes beyond mere orientation the sense that it is also provides the ‘evidential’ weight to be assigned to certain properties of a stimulus in accordance with how diagnostic they are in categorization.³⁰

Third, attentional weighting is part of a broader process of perceptual learning that involves coming to represent (and weight in attention) the relations between different characteristic properties of the object of interest. In stimulus imprinting, the visual system develops specialized detectors for stimuli or parts of stimuli (Goldstone 1998). Topological imprinting is a kind of stimulus imprinting whereby a network of detectors is created and organized spatially, allowing for the detection and encoding of characteristic relations. For this reason it is considered to be able to detect more abstract properties of stimuli, as it allows for the

³⁰ A further reason to deny the judgment hypothesis is that some of the evidence for attentional weighting is a phenomenon known as ‘categorical perception’ (see section 4 below). This stable weighting of different features for categorization has an enduring effect on perceptual experience.

encoding of not just of the presence or absence of a feature, but rather the characteristic relations between them (sometimes called ‘second-order relations’), where this may be a more or less variable relation, depending on the object. There is not that much variation in how noses, eyes and mouths are spatially organized and proportioned, whereas the various properties of dogs (tails, noses, paws, etc.) may exhibit higher variation due to the large differences between breeds.

Evidence for both stable attentional weighting for categorization and topographical imprinting comes in part from the fact that experts often process the objects of their expertise in a holistic manner. In such cases, experts (but not novices) with a given category are unable to ignore properties that are diagnostic for categorization but irrelevant for a transitory task. For example, the composite task uses a stimulus that is a composite of two different exemplars (Gauthier et al. 2003; A. W. Young, Hellawell, and Hay 2013). The top of Arnold Schwarzenegger’s face might be combined with the bottom of Danny DeVito’s face. Subjects are either shown two composites at a time, or one after the other. The task is to attend only to (say) the top half of both composites in order to judge whether both top halves are identical or different. Trials are either congruent or incongruent. In incongruent trials, one part changes and the other remains the same. A subject might be shown the top of Schwarzenegger’s face combined with DeVito’s face, along with another complete image (both top and bottom halves) of Schwarzenegger’s face. In congruent trials, both parts change or remain the same. A subject might be shown the Schwarzenegger/DeVito composite twice, or alongside a different composite face, say of Robin Williams and Will Ferrell.³¹

³¹ The composite task has also been performed with car novices and experts (Gauthier et al. 2003), for Chinese character novices and experts (Wong et al 2012), and also with people trained to recognize lab-manufactured objects called ‘greebles’ (Gauthier et al. 1998) and ‘ziggerins’ (A. C.-N. Wong, Palmeri, and Gauthier 2009).

Evidence for topographical imprinting and stable attentional weighting for categorization comes from expert performance in the incongruent trials, where either their accuracy is lower or their response times are slower than in the congruent trials. A common way of interpreting this evidence is that the properties relevant to facial recognition (or to recognition of other objects of expertise) are grouped together as units, with the attentional weights of these properties relatively stable across contexts, and the relations between the various properties important to recognition. Though subjects are instructed to pay attention to only one part of the stimulus in order to perform the task, they are able to do this only with difficulty because this requires a different, transitory attentional pattern to properties, and a disregard for the second-order relations between some properties (we cannot rely on the distance between nose, mouth, and eyes to determine whether there is a match).

This then constitutes a third way that the conception of perceptual expertise as a mere learned pattern of attention is mistaken. The end result of perceptual learning for expertise is a perceptual category, understood as an interconnected network of weighted property detectors, where the relations between these properties may also be weighted.³² The expert not only learns to weight certain properties more or less heavily in attention, but also to represent new relations between individual properties, and to weight these properties according to their diagnosticity for categorization. Again, this demonstrates that the difference between novice and expert cannot be explained in purely attentional terms: experts represent relations between properties that novices cannot.

³² The resulting construct used for perceptual categorization may strike some readers as sounding much like a prototype. There are indeed prototype models of perceptual categorization, and this is the model I favour here as well. However, exemplar models and other proposed computational models, also make use of differential attentional weighting of diagnostic properties (Palmeri and Cottrell 2010).

While these points count against simpler forms of the judgment hypothesis, such as that proposed by Pylyshyn, they cannot yet rule out the more sophisticated proposal put forth by Burge. On his elaboration, what the perceptual category detects can be explained as a mere assemblage of low-level properties. If this is the case, then the judgment hypothesis is compatible with perceptual learning for expertise described thus far.

However, Burge's perceptual category must nevertheless be complex enough to track the high-level property in question, or else it is no longer a useful construct for explaining the skills of the expert. The expert can only categorize objects reliably and rapidly if she already has all the perceptual evidence assembled. Once this is granted, then – as I will argue here – this renders the post-perceptual judgment of high-level properties redundant. The perceptual category alone possesses sufficient structure so as to be able to represent high-level properties directly in perceptual experience.

The simplest kind of mere assemblage will be a bare list of low-level properties. This clearly cannot track high-level properties. Take the example of the high level property 'dog.' Dogs vary greatly in their low level properties, such as the colour of their fur and their size. A bare list would provide all these alternatives (brown, black, red, etc.), but would need to be supplemented by something like a measure of typicality or frequency indicating for each feature how diagnostic it is for category membership. This is because there must be some means of indicating that not all items on the list need to be detected in order for categorization to be made. A dog needn't (and likely isn't) all of these colours at once. A bare list also provides no guidance as to how the properties on the list are related to one another, and so must be supplemented by the inclusion of second-order relations. It would not do, for instance, to detect fur in one part of

the visual field and a tail in another and – despite there being a gap of sky in between – categorize this as a dog.

The perceptual category must then be able to track the pattern of interest to the expert. Patterns can be understood along lines proposed by (Dennett 1991). Drawing on Gregory Chaitin's definition of mathematical randomness, Dennett proposes that: "A series (of dots or numbers or whatever) is random if and only if the information required to describe (transmit) the series accurately is incompressible: nothing shorter than the verbatim bit map will preserve the series. Then a series is not random – has a pattern – if and only if there is some more efficient way of describing it" (1991, 32). On this definition, patterns involve the compression of information: elements that are not shared by members of the series or category can be discarded, and those elements that are shared and relevant to membership are preserved as constitutive of the pattern.

Instances of high-level properties are instantiations of a given pattern. While any given instantiation may be explained in terms of a mere assemblage of low-level properties, the pattern cannot be so explained. For example, an individual dog might be described entirely in terms of low-level properties such as colour, size, shape, and the individual relations between them. In this sense, we might suppose that even second-order relations are describable as simple relations along lines that sparse content theorists would endorse. Such theorists have no issue with positing that we represent simple relations between properties in perceptual experience, such as that the sphere is to the left of the rectangle.

However, the pattern corresponding to 'dog' cannot be explained this way. The perceptual category in the business of detecting such pattern instances will involve a selection and compression of information that is no longer describable in terms of specific low-level

properties. The pattern is not reducible to any particular instance. Indeed, it cannot be, or else it fails to capture the other instances it attempts to capture. This is particularly evident in the case of the second-order relations. The second-order relations between relevant properties encoded in perceptual categories include not just a simple spatial relation, but also the lower and upper boundaries of those relations, along with their typicality. This information must be encoded in order to handle the inevitable variation across instances: even though a dachshund's torso is much longer than is typical, we still recognize it as a dog. The judgment hypothesis cannot explain the pattern of detection without reference to the encoding of more complex second-order relations.

Once the inclusion of these items is posited, then post-perceptual judgment in order to ascribe a high-level property is arguably rendered redundant. The systematic encoding of second-order relations between these low-level properties, as well as the evidential weighting of such properties is sufficient to diagnose in perception membership in a high-level category for most exemplars of that category. This renders post-perceptual judgment unnecessary for categorization, because the work can be carried out in perception.

Consider an analogy: in order to determine whether it is raining, we might adopt a fairly stable policy for how to weight the evidence we receive. Perhaps direct visual evidence such as seeing falling droplets will be given the highest weighting, followed by testimony, followed by seeing someone holding an umbrella, and so on. Once we have a policy as to what evidential weight to assign each of these, and have adopted a threshold for when we will form a belief that it is raining, then we do not need to make further inferences to judge that it's raining. Being sensitive to and weighing the evidence according to the pre-established weights is all we need to get to the judgment.

The idea is that it is the same for attentional weighting and stimulus/topographical imprinting: stimulus imprinting serves to be able to detect the relevant evidence and group it together – it ensures we are sensitive to the relevant properties and their relations. Attentional weighting serves to set how much weight to give each bit of evidence. The perceptual category as a whole sets the threshold for categorization.³³ On this view then, the process of attentional weighting does not support the judgment hypothesis. The abilities of experts – when understood as resulting from perceptual learning for expertise – should be understood as allowing for the representation of new high-level properties in perceptual experience. While this falls short of establishing that perceptual kinds are natural kinds, it nevertheless provides compelling evidence that they are high-level properties.

Here I suspect that some readers will have in mind the following objection: even if we grant that perceptual kinds are high-level properties, why think that they are represented in perceptual experience? Why not hold that the categorization process occurs post-perceptually instead, and so ascription of perceptual kinds is a post-perceptual judgment after all? After all, where to draw the division between perception and cognition is contested, and so we might think that the process described here ought to be labeled cognitive, given its sophistication. I answer this objection in the next section.

³³ While it is true that in some cases we will need instead to rely on post-perceptual judgment – such as in suboptimal perceptual conditions, or when confronted with highly atypical exemplars (penguins, echidnas), or in cases of perceptual illusions, these will be the exception rather than the rule. In such cases, perceptual categorization fails. Also, while how much weight to assign a given source of evidence is often context-sensitive, categorization will be a relatively stable task that cuts across contexts, and when there are relevant contextual variations, these will be relatively stable regularities as well, and so also open to learning. For example, we may learn to downgrade the attentional weight accorded to auditory evidence in very noisy contexts, or to visual evidence in low lighting contexts.

3.3 The learned properties are represented in perceptual experience

The common mechanism argument involves appealing to a characteristic feature or effect of the perceptual processing of low-level properties. If this effect is also present in the processing of a given high-level property, then this is evidence that the high-level property is also represented in perceptual experience. Such an argument is implicit in the many psychological studies discussed below that look at so called ‘high-level perception,’ and it has become increasingly common for philosophers to appeal to some version of this argument in order to adjudicate the content debate (Bayne 2009; 2016; Block 2014; 2016; Briscoe 2015; Brogaard 2018; Fish 2013; Mandelbaum 2018; Stokes 2020).

An adapted version of the common mechanism argument can be applied to perceptual learning for categorization, because many studies of perceptual expertise involve comparing experts and novices in a given domain, or comparing a given population prior to and subsequent to training. Evidence for our abilities to make basic-level categorizations is also relevant if we assume that these are learned, albeit much earlier in most of our lives. Here the relevant contrasting population will mostly consist of infants and children.³⁴ The argument again relies on isolating a feature or effect of the perceptual processing of low-level properties, and then pointing to evidence that our attribution of a certain high-level property exhibits this effect after, but not prior to learning.

³⁴ In what follows I will not review the developmental evidence for categorization. The literature is too vast, and there are too many subtleties in interpreting the evidence in order to do it justice here (Rakison and Yermolayeva 2010). If people are unconvinced that our abilities to make basic-level categorizations are learned, then the evidence reviewed below can be understood as supporting the claim that such innate abilities are best understood as perceptual.

The argument can be summarized as follows:

1. Prior to perceptual learning attribution of high-level property F does not exhibit effect G
2. After perceptual learning attribution of high-level property F does exhibit effect G
3. The best explanation for (2) is that high-level property F has come to be represented in perception
4. Therefore, high-level property F is probably represented in perception

The sorts of effects commonly appealed to in the common mechanism argument include speed, automaticity, encapsulation, and adaptation. Here I will provide a brief review of evidence from these categories. Much of this evidence comes from the study of face expertise, and so a few remarks are in order. First, it is important to distinguish between distinct abilities we possess with respect to faces. We can categorize a face as a face (basic-level category), as a face of a particular race, approximate age or gender, as expressing a certain emotion (all subordinate categories), or we can recognize it as belonging to a particular individual. Care should therefore be taken when reviewing the evidence to note which ability is under consideration as perceptual. While evidence for one ability being perceptual may to some degree count in favour of the other abilities being perceptual, the ideal is to have direct evidence for each ability related to faces. Here I admittedly fall short of this ideal, in part due to space considerations, and in part due to lack of relevant empirical data for some of these categories.

Second, while there is controversy over whether these face-related abilities are learned or innate (Sugita 2009; Arcaro et al. 2017) there is nevertheless much evidence to suggest that there is room for learning either way. We can get better at recognizing faces (Limbach et al. 2018),

emotional expressions (Pollak et al. 2009; Pollak and Sinha 2002), individuals (Itz et al. 2014; Limbach et al. 2018) and individuals of other races (Delaney et al. 2011; Sangrigoli et al. 2005; Markant and Scott 2018). So I take it that even if we are hard-wired to either detect or learn about faces, this does not discount the role of perceptual learning in our gaining expertise. This addresses a potential worry that research on faces cannot be used in the version of the common mechanism argument I have proposed here, because our expertise in this case is not learned, and so there is no relevant novice contrast class.

Finally, research on face expertise is also important because it has become increasingly apparent that other objects of acquired expertise – such as cars, birds, fingerprints, and the like – recruit the same neural substrates as are used to process faces once expertise has been achieved. Such non-face objects are processed in the same time course as faces, in the same localized brain region, and cause interference in face processing when presented simultaneously with faces to people who are experts with the non-face objects (Curby and Rossion 2010; Gauthier et al. 2000). The evidence for representing faces in perception reviewed below thus carries some weight in supporting the claim that other objects of expertise can come to be represented in perception, even if one believes that we are born face experts. The sorts of effects exhibited by faces reviewed below provide good reason to hold that they are represented in perception, and the shared time course and neural substrates of face processing then count in favour of the acquired non-face expertise also being perceptual. This evidence thus puts pressure on those sympathetic to the claim that we represent faces in perceptual experience, but who deny that we perceptually represent any culturally-learned items (Block 2014; Burge 2010; 2014)).

3.3.1 Speed

Perceptual processing takes place quite rapidly. Neural responses to individual low-level properties such as flashing lights occurs as quickly as 30-50 ms (Wilson et al. 1983).

Conjunctions of low-level properties are processed more slowly, however (Bodelón, Fallah, and Reynolds 2007), and feature binding and object segregation have been hypothesized to occur during what is known as ‘local recurrent processing,’ approximately 100-150 ms after stimulus onset (Lamme and Roelfsema 2000). If high-level properties were processed within this time window, then this would count against the judgment hypothesis: we would not likely have time to draw on cognitive resources in order to make a categorization.

The most compelling evidence for the perception of high-level properties comes from the speed of some basic and superordinate-level categorizations. Using single-cell recordings from implanted microelectrodes in human epileptic patients, category-specific neural firing has been found to occur within 100 ms for human faces, fruit, animals and chairs (H. Liu et al. 2009). Using behavioural measures (Kirchner and Thorpe 2006) estimate that categorization of scenes containing animals occurs within 95-100 ms, and (Thorpe, Fize, and Marlot 1996) found ERP responses for categorization of scenes as containing an animal to occur by 150 ms.

The categorizations that perceptual experts are able to make are also much faster than those of novices. The measured brain response, or ERP, for faces has been found to occur approximately 130-170 ms after a person has been exposed to an image of a face, and is

commonly referred to as the N170 (Bentin et al. 1996; Rousselet et al. 2008).³⁵ It is thought to pertain to the categorization of faces as faces, rather than identification of a specific individual (Allison et al. 1999). A heightened N170 response has also been found for bird and dog experts viewing the objects of their expertise, but not when viewing objects outside of their expertise (Tanaka and Curran 2001). Training novices in bird categorization also produced a heightened N170 response (Scott et al. 2006), with the same findings extending to novices trained to categorize car varieties (Scott et al. 2008). Finally, novices that have been trained to categorize lab-manufactured objects called ‘greebles’ (which look kind of like gremlins) also come to exhibit an N170 response when presented with images of greebles (Rossion et al. 2002).

ERP evidence of gender and age categorization of faces has been found to occur at 145-185ms after stimulus onset, though this is thought to be independent of the N170 (Mouchetant-Rostaing, Giard, Bentin, et al. 2000; Mouchetant-Rostaing and Giard 2003).

There is also a later ERP component, called the N250 (occurring at approximately 250 ms after stimulus onset), that is thought to correspond to recognizing faces as belonging to specific individuals (Tanaka et al. 2006), as well as to being able to make distinctions at the subordinate level of categorization rather than at the basic level (German shepherd vs. lhasa apso) (Scott et al. 2008). However, this clearly falls outside the time window of local recurrent processing for low-level properties. Does this mean that the processing is cognitive?

Determining the cutoff speed for perceptual processing is controversial – it requires taking a stand on what brain regions count as ‘cognitive’ or ‘perceptual,’ something that itself relies on

³⁵ There is some controversy, however, over the N170 as the earliest locus of detection of faces. (Seeck et al. 1997) found differential ERP responses to novel vs. familiar faces within 50-70 milliseconds. See (Vanrullen and Thorpe 2001) for some discussion of these results.

background theory. So while these last findings do not clearly support the perceptual expertise hypothesis, neither do they count against it.

3.3.2 Automaticity and encapsulation

Automaticity concerns whether or not an agent can prevent the processing of a class of stimuli. If the processing is automatic then, upon exposure to a stimulus, the system processes the information whether the agent wants to or not (Mandelbaum 2015).³⁶ Low-level properties are thought to be processed automatically in perception: the only way the agent can interfere with this process is by cutting off the source of stimulus input, such as shutting her eyes. A system or process is encapsulated if its operation is sensitive only to information contained within the system (Fodor 1983; Pylyshyn 1999). That at least early perceptual processing is encapsulated from cognition is supported by the existence of perceptual illusions. For example, no matter that we believe that the lines of the Müller-Lyer illusion are of the same length, they continue to appear to us as if they were different lengths.

There is evidence that experts process the objects of their expertise automatically as well, and that this process is encapsulated. One source of evidence for subordinate categories is that prior to expertise objects are processed in a parts-based fashion, and after expertise objects are processed holistically, or configurally. I discussed the composite task in detail in section two as evidence for the claim that perceptual learning for categorization involves stable attentional weighting and the grouping together of features and the relations between them. However, the

³⁶ While (Mandelbaum 2015) also discusses automaticity in terms of processes that cannot be interrupted once initiated, lab techniques such as backward masking show that it is possible to interrupt perceptual processing.

composite task also supports the claim that processing of objects of expertise is automatic and encapsulated. Despite the attempts of experts to ignore the task-irrelevant portions of faces (or cars, or greebles), they are not able to do so. They cannot prevent such configural processing, and their beliefs concerning the way to successfully perform the task do not have an effect, except after the initial holistic processing.³⁷

This sort of evidence counts against the judgment hypothesis. Since post-perceptual judgments are cognitive, they should be sensitive, at least to some degree, to our other beliefs and desires – they should not be encapsulated, and so should be possible to disrupt. The sort of sensitivity characteristic of post-perceptual judgments is well-documented by (Firestone and Scholl 2016) who discuss many experiments that illustrate that supposed effects of cognition on perception are instead explained in terms of an effect on post-perceptual judgments. For example, while one experiment concluded that throwing a heavy ball made the target look farther away on the basis of subject estimates of the distance, a follow-up study that explicitly instructed subjects to make the distance estimate solely on the basis of visual information found no effect due to the heaviness of the ball (2015, p.9). Firestone and Scholl highlight the susceptibility of post-perceptual judgments to experimenter demands and response bias: while another study concluded that wearing a heavy backpack caused a hill to appear steeper to subjects based on their estimate of the slant, a follow-up that provided subjects with a (false) story about why they had to wear a backpack saw the original effect completely disappear (2015, p.10).

The wealth of examples discussed in (Firestone and Scholl 2016) provides compelling reason to think that post-perceptual judgments are not automatic and encapsulated in the way that

³⁷ For a discussion of how automaticity may be measured by the (lack of) effects of conscious expectations on perceptual categorization, see (Tanaka and Curran 2001, 45).

perceptual categorization is, and so the evidence reviewed above counts against the judgment hypothesis.

3.3.3 Adaptation effects

Adaptations are rapid and temporary adjustments in the sensitivity of our sensory systems that result in changes in perceptual awareness (Webster 2012; 2015). Many low-level properties are subject to adaptation effects. For example, adaptation to a scene coloured one way causes a black and white version of the same scene to appear coloured in an opposite fashion (the original red areas are perceived as green, and so on).³⁸

Though adaptation is not typically understood as perceptual learning because the changes tend to be short-lived rather than persistent over time, it is nevertheless relevant to the common mechanism argument: if high-level properties fail to exhibit adaptation effects prior to learning but then exhibit them after learning, then this is evidence in favour of their representation in perception (McGovern, Roach, and Webb 2012; though see the following for discussion of interaction between perceptual learning and adaptation Kompaniez-Dunigan et al. 2015; Kompaniez et al. 2013). Or, if high-level properties exhibit adaptation effects, and adaptation effects are proprietary to the perceptual system, then this is evidence in favour of their being perceptual.

Faces are a well-studied source of adaptation effects (Webster and Macleod 2011). For example, exposure to a horizontally distended face causes the original undistorted face to appear

³⁸ For low-level property adaptations, see for example (Clifford 2002; Durgin and Proffitt 1996; Suzuki and Cavanagh 1998).

horizontally constricted (Webster and Maclin 1999; Susilo, McKone, and Edwards 2010). The recognition of individual faces can also be manipulated through adaptation, with adaptation to an individual's 'anti-face' causing an averaged face to appear more like the original individual (Leopold et al. 2001). Emotional expressions are also subject to adaptation (Ellamil, Susskind, and Anderson 2008; Fox and Barton 2007). For example, fixating on a fearful face will cause a subsequently presented ambiguous face to appear angry, and fixating on an angry face will cause the ambiguous face to appear fearful (Butler et al. 2008). Finally, adaptation effects have also been observed for facial age (Schweinberger et al. 2010), as well as gender and ethnicity (Webster et al. 2004).

Beyond faces, adaptation has been found for gender, using gendered bodies without showing the face (Palumbo, Laeng, and Tommasi 2013) and global scene properties such as 'openness' and 'naturalness' (Greene and Oliva 2010). Adaptation effects have not been studied for other areas of perceptual expertise, however.

In summary, the combined evidence here provides support for the perceptual expertise hypothesis over the judgment hypothesis. While this support is not decisive, and further empirical evidence may strengthen or weaken this support, this is just as expected for an abductive argument. An additional source of important evidence that one might appeal to on the common mechanism argument is categorical perception – it is an effect present in the processing of some low-level properties that is also found for some high-level properties. However, I review it separately in the next section because it also serves as evidence that high-level properties are represented in the phenomenal character of perceptual experience.

3.4 There is a change in the phenomenal character of perceptual experience due to high-level properties

One might think that, from the arguments made in the previous sections, I have already adequately defended the claim that we represent high-level properties in perceptual experience, where perceptual experience is understood as ‘what it is like’ to have that perceptual experience – its phenomenal character. However, whether this is the case depends on a further question as to how the contents of perceptual experience and the phenomenal character of perceptual experience are related. On one theory – called representationalism or intentionalism – the representational content of perceptual experience determine the phenomenal character of experience (Egan 2012). On the strongest version, such determination is exhaustive – when there is a change in what our perceptual experience represents, then there must be a change in what our perceptual experience is like (Tye 2002). In such cases, the link between the contents of perceptual experience and the phenomenal character of perceptual experience is straightforward – once it is established that we represent a high-level property in perceptual experience, then it is a given that this will alter what this perceptual experience is like for us. If you already accept this sort of view then you can read the following section simply as a further application of the common mechanism argument.

However, one might hold the view that the representational content of perceptual experience does not always impact phenomenal character – the contents of perceptual experience can outstrip what our experiences are like. If this is the case, then it may be that while we represent high-level properties in perception, these properties have no impact on the phenomenal character of our perceptual experience. (Prinz 2013) argues for such a view, labeling this position

‘phenomenal modesty,’ in contrast to ‘phenomenal richness,’ the view that high-level properties represented in perceptual experience do impact the phenomenal character of perceptual experience. (Chudnoff 2018) also argues that while perceptual learning can allow us to represent high-level properties in perceptual experience, they nevertheless lack ‘presentational phenomenology,’ where this can be roughly understood in terms of phenomenal character (for further discussion see Brogaard and Gatzia 2018).

This section of the chapter provides an argument for phenomenal richness that does not depend on adopting a particular theory of how the representational content of perceptual experience and the phenomenal character of perceptual experience are related (from here on I will refer to the phenomenal character of perceptual experience simply as ‘phenomenal character’ for brevity). My strategy is to appeal to a perceptual effect that was not discussed in the previous section, though it is also a candidate for the common mechanism argument elaborated above: categorical perception. In the case of low-level properties, categorical perception demonstrates that such properties impact phenomenal character. Insofar as these effects also occur with respect to high-level properties we should also accept that the impact of high-level properties is to phenomenal character.

Categorical perception is usually described as the perceptual shifting of the world according to some of the categories we possess. For example, (Goldstone and Hendrickson 2010) write that “our perceptions are warped such that differences between objects that belong in different categories are accentuated, and differences between objects that fall into the same category are deemphasized” (p. 69). Things from the same category look more similar to us, and things from different categories look more different. However, this way of describing the

phenomenon may be construed as question begging – why should we think that categorical perception is the effect of a category on phenomenal character?

One reason for doing so is that the effect is observable in our own everyday perceptual experience of low-level properties represented in perception. For example, when we see a rainbow we perceive relatively discrete bands of colour rather than a mere continuous gradation. This is thought to be an effect of categorical perception – despite the fact that there is a continuous spectrum of wavelengths, our visual systems parse the spectrum into discontinuous categories (Bornstein and Korda 1984). Our colour categories cause us to perceive certain wavelengths as more similar to one another, and others as more different, despite the constant physical differences. The effect is striking, much in the way visual illusions are striking once one realizes that it is an illusion – one can easily observe categorical perception of low-level properties for oneself. This is true not just of colour but of the other low-level properties such as phonemes, which are discussed below.

Discrimination testing is a way of bypassing reliance on introspection and first person reports, and so offers additional support for the claim that categories have an effect on phenomenal character. A consequence of categorical perception is that discriminating between things that are of the same category is more difficult, and discriminating between two objects that are of different categories is easier, even when physical differences between the objects are controlled for (Harnad 1987; 2003). Therefore, testing for categorical perception usually involves creating a series of stimuli that vary equally along a continuum that are then used in a discrimination task. The discrimination task often takes the form of an ‘ABX’ task, where subjects are typically shown two stimuli next to each other in the continuum, followed immediately by a third stimulus that is identical to one of the first two stimuli. The task is then to

say whether X is identical to A or B. Either accuracy, speed, or both are used as measures of performance, and when these differ systematically, depending on whether the stimuli are from the same or different categories, then this is evidence of categorical perception. (An identification or categorization task is also usually given after the discrimination task to determine the subjective category boundaries.) For example, (Liberman et al. 1957) constructed a continuum of syllable sounds from /be/ to /ge/, with physically equal transformations between each sound. Subjects were better able to perform a discrimination task when the phonemes were from different categories than when they were from the same category. They also consistently reported an experience not of a smooth transformation, but rather the experience of rather abrupt sound transformations, from /be/ to /de/ to /ge/. The ability to discriminate between two objects is reasonably thought to correspond to our perceptual experience of them as the same or different, and so we should take this testing as a proxy for phenomenal character.

Because the discrimination task involves working memory, there is always the worry that what is going on in such cases is describable as follows: when we are shown the two initial AB stimuli, if they are from the same category we ‘hold’ them both in working memory as such, and if they are from different categories then we hold them both in working memory as different category members. Then when we are shown the subsequent X stimulus, of course we will be more accurate or faster at discrimination in cases where the initial A-B pair are from different categories, because we have encoded them that way. The change in ability can be explained in terms of learned category, but the test does not measure what we think it measures – a change in phenomenal character. Instead it reflects a mnemonic strategy that is gained with category learning.

To rule out this possibility, sometimes the discrimination task is conducted with two simultaneously presented stimuli, though this is not feasible in the case of auditory stimuli (Harnad 1987). The basic idea is the same – either the two stimuli presented are from the same or different categories, and the subject must judge whether the two stimuli are the same or different. If the subject performs better in situations where the stimuli are from different categories, then this is evidence for categorical perception. Goldstone and Hendrickson (2010) review additional, though mixed, evidence for an early perceptual locus of processing.

Once we allow that the categorical perception of low-level properties has an impact on phenomenal character, then we should allow that categorical perception of high-level properties – insofar as it exists – has an impact on the phenomenal character of perceptual experience. Few would challenge the claim that the phenomenal character of our perceptual experience of low-level properties is in fact warped or shifted in the ways described. While one might mount an argument that the shift in phenomenal character is not due to the categories we possess but some other unspecified factor, this argument is undermined by cases of perceptual category learning – in such cases the shift in discrimination abilities prior to and after acquiring the relevant category is best explained by acquisition of the category itself.

A nice example of this is musical pitch. In a comparison of the pitch discrimination abilities of expert and novice musicians, a categorical perception effect was found at the semitone boundaries only for experts (E. M. Burns and Ward 1978). Other cases where categorical perception has been observed as the direct result of learning include people trained to recognize artificial speech-like categories (Lane 1965), English speakers trained to recognize phonemes that appear in foreign languages but do not appear in the English language (Pisoni et

al. 1982), and people trained to recognize squares either in terms of brightness or size categories (Goldstone 1994).

These sorts of cases support the claim that it is the learned category that is responsible for categorical perception: prior to category learning there is no categorical perception effect, and after category learning there is. There are also additional considerations that favour the claim, such as the ability of mathematical and computational models of categorical perception to predict the observed effects on performance (Goldstone and Hendrickson 2010). In such models, it is some aspect of the way the stimuli have been categorized that leads to the effects.

The next step of the argument is then to support the claim that categorical perception does occur for high-level properties. It is more challenging to test for categorical perception of high-level properties in part because the stimuli are relatively dimensionally complex as compared to simple stimuli such as colours and phonemes, and so it is hard to create stimuli that vary continuously along a spectrum. Nevertheless, there is some evidence that categorical perception is present for everyday objects such as bottles, vases, and bells – objects that can be morphed via computer into each other with relative ease and continuity (F. N. Newell and Bülthoff 2002). Categorical perception is also observed in people who have been trained to categorize laboratory-created complex objects somewhat resembling biological cells (Livingston, Andrews, and Harnad 1998). Finally, categorical perception effects have been found for human facial identity (Rotshtein et al. 2005), emotional expressions (Calder et al. 1996; Etcoff and Magee 1992; Fugate 2013), gender (Campanella, Chrysochoos, and Bruyer 2001) and race (Levin and Angelone 2002).

While limitations to testing make it difficult to test a wide variety of high-level properties, the cases of categorical perception discussed here nevertheless provide proof of concept – if we

accept that the high-level property in these cases makes a difference to the phenomenal character of perceptual experience, then we should accept that other high-level properties represented in perception also make a difference to the phenomenal character of perceptual experience.

3.5 Conclusion

In this chapter I have provided an argument for the perceptual expertise thesis: through perceptual learning we can come to represent high-level properties in perceptual experience. My argument involves three steps. First, I draw on the mechanics of perceptual learning to argue that the properties really are high-level. Attentional weighting and stimulus imprinting are two processes by which the perceptual learning of categories occurs. Stimulus imprinting groups together the features that are characteristic of a given category (along with the relations between those features), and attentional weighting assigns weights to these features in accordance with how diagnostic they are for this category. The result is a stable perceptual category that is capable of attributing high-level properties to objects. Positing a separate post-perceptual categorization process is superfluous.

Second, I appeal to a variant of what I term the common mechanism argument in order to argue that these properties are represented in perceptual experience. The common mechanism argument begins from common ground between rich and sparse content theorists by appealing to effects of the perceptual processing of low-level properties. If such effects can also be found in the perceptual processing of high-level properties, then this is evidence in favour of the claim that those high-level properties are also represented in perceptual experience. Perceptual learning provides additional a way of running a variant of this argument because it provides a contrast

between the way in which novices and experts process the high-level properties of expertise. Experts process such properties more quickly and automatically than novices, and adaptation effects are exhibited for some such properties. This change in processing is best explained by positing that the properties come to be represented in perceptual experience.

Third, I argue that representing high-level properties in perception alters the phenomenal character of perceptual experience – the perceptual categories we possess determine in part what it is like to perceive the world. I appeal to categorical perception as supporting evidence. Categorical perception is an observable psychological shift in our perceptual experience due either to learned or innate perceptual categories, where objects that are members of the same perceptual category appear more similar, and those of different categories appear less similar. This is sufficient to demonstrate that high-level properties can alter phenomenal character.

Chapter 4: Learning to See Beauty: A novel account of aesthetic perception

Summary: Do we perceive the beauty of paintings, symphonies and sculptures, or are our judgments of beauty non-perceptual? While the perceptual view may seem intuitive, it has received little in the way of explicit defence. It also faces multiple hurdles, including accounting for the role of aesthetic training. Unlike our perception of shapes, colours and the like, training is often involved in our coming to make aesthetic judgments, particularly with respect to artworks. This puts pressure on the perceptual view because it suggests that we rely on knowledge of the artworks, rather than perception to ascribe aesthetic properties. In this chapter I provide an argument in favour of aesthetic perception by way of developing a positive view of how we can come to represent aesthetic properties in perceptual experience. This view builds on and develops themes from Kendall Walton's 1970 paper 'Categories of Art'. As such, I call the view 'Waltonian perceptualism.' Amongst other virtues, this view explains how aesthetic properties can be both learned via training and genuinely perceptual.

4.1 Introduction

Many philosophers have held some version of the view that we either perceptually experience, or experience in some broader sense, aesthetic properties (Iseminger 2004; Wollheim 1970; Zangwill 1998; Fudge 2005; Hopkins 2005; 2006; Lopes 2014a; 2016; Tormey 1973; Irvin 2008b; Levinson 2005; Pettit 1983; Lamarque 2010; Stokes 2014; 2018a; De Clercq 2002). The strongest version of this view, and the view that I will primarily focus on here, is that we represent at least some aesthetic properties in perceptual experience.³⁹

Whether or not this view is correct matters for our understanding of the epistemology of aesthetic beliefs. For example, if we perceive aesthetic properties, then this suggests that the justification for aesthetic beliefs may be epistemically immediate (Pryor 2005; Alston 1983; Dorsch 2013). That is, our aesthetic beliefs may not depend upon other beliefs for their justification if they are formed on the basis of perceptual experience alone.

Aesthetic perception would also reveal something about the nature of perceptual experience more generally. Whereas low-level properties such as shape, size, motion and colour are universally acknowledged to form part of the contents of perception, it is controversial whether we perceive high-level properties such as artifactual or natural kind properties, aesthetic properties and moral properties (Stokes 2018a; 2014; Siegel 2010; Brogaard 2013; Siewert 1998; Block 2014; Burge 2014; 2010). Insofar as aesthetic properties are high-level properties, then their being perceived would count in favour of the view that we represent high-level properties in

³⁹ While I do not have the space to discuss it here, aesthetic perception is related to, but distinct from aesthetic empiricism (Gregory Currie 1989).

perceptual experience. Aesthetics is thus well poised to make a contribution to the philosophy of perception.

Finally, the answer has implications for how we understand aesthetic experience. The debate on this subject has been largely centered on whether or not there is such thing as distinctively aesthetic experience, and if so, what characterizes it (Carroll 2002; Irvin 2008a; Lopes 2014a ch. 9). If we perceive aesthetic properties then this should reveal something about the nature of aesthetic experience: that it involves the perceptual experience of aesthetic properties. Whether this constitutes the entirety of such experience is another matter left undecided by a verdict on the perceptual question. However, any tenable position on aesthetic experience would need to include a perceptual component.

In this chapter I provide a positive argument for aesthetic perception (see also Ransom 2020). In section 4.2, I review reasons for and against adopting a perceptual view. While none of these are decisive, the reasons against adopting a perceptual view nevertheless provide us with desiderata – any perceptual view should be able to address these issues in a plausible manner. In section 4.3 I review the prospects for a direct argument in favour of the perceptual view, and find extant attempts are not convincing. The best prospect for the view is a plausible elaboration of how aesthetic perception actually occurs. In section 4.4 I present and critically evaluate such an attempt: that the way we come to perceive aesthetic properties is via cognitive permeation. In section 4.5 I provide a full account of how some properties are represented in perceptual experience: perceptual learning explains how we can come to represent new categories of art and objects of aesthetic appreciation in perception, and the relative fluency with which such objects are categorized explains how we can represent aesthetic properties in perceptual experience. An account of what it is to be an aesthetic expert provides us with the means of saying when those

perceptual experiences are veridical, and the objects can truly be said to have those aesthetic properties. Finally, in section 4.6 I address several objections and discuss how my account can accommodate the challenges to aesthetic perception raised in section 4.2.

4.2 Considerations for and against aesthetic perception

4.2.1 Language

To argue that aesthetic properties are perceived, we might first turn to the way we use language (Anna Bergqvist and Cowan 2018; Logue 2018; Stokes 2018a). The view that we perceive aesthetic properties seems to be reflected in our everyday talk. We speak of *seeing* the beauty of a painting, *hearing* the melancholy tone in a piece of music, and *tasting* the fineness of a wine. However, this sort of consideration provides little to no support in favour of the rich content view due to the fact that we also use perceptual terms quite extensively in a loose or metaphorical sense. We speak of *seeing* the problem when it would be more apt to speak of understanding it, or the sweet *taste* of success when the success in question doesn't literally involve eating or drinking something. Deciding whether or not aesthetic talk is literal or metaphorical will require further consideration.

4.2.2 Psychological immediacy

When we stand in front of a painting, go to a movie or read a book, we seldom form our judgments by explicitly considering reasons for or against judging the artwork to be powerful,

beautiful, or lacking. Another way of saying this is that our ascription of aesthetic properties are psychologically immediate – we do not arrive at such ascriptions by way of conscious inference. If they are not consciously inferred, then this would seem to provide some support for the claim that we call something beautiful because we perceive it to be such.

However, psychological immediacy is compatible with another explanation: that our ascription of aesthetic properties takes place via unconscious inference.⁴⁰ While we do not consciously weigh reasons, there is an unconscious process that causes us to judge or infer, rather than perceive, aesthetic properties. Call this position *inferentialism*. On this view, we perceive only low-level properties (or perhaps also non-aesthetic high-level properties) and then infer the presence of aesthetic properties (Dorsch 2013; Logue 2018; Cavedon-Taylor 2017). This inference need not take place consciously – it may be unconscious or implicit, in much the same way that we may come to form the belief that it is raining outside upon seeing someone walk into a building with a wet umbrella without having consciously gone through an inferential process. Psychological immediacy offers little to no support for the perceptual view.

4.2.3 Aesthetic diversity

The wide range of things that we may judge beautiful includes people, sunsets, paintings, and music, amongst other objects. The sheer diversity of the category of ‘beautiful things’ poses a *prima facie* challenge to aesthetic perception because it weighs against the idea that there are any common perceptible features shared by all or most members of the category. This is unlike

⁴⁰ There are other alternatives to the perceptual view that I do not consider here, such as the view that aesthetic properties are represented in non-perceptual experiences via the emotions (Anna Bergqvist and Cowan 2018). I take the positive argument of this chapter to equally pose a challenge to these alternatives, however.

the case of many other high-level properties argued to be represented in perceptual experience. For example, pileated woodpeckers share their colour, size, and patterning. While dogs are a more perceptually diverse group, there are still perceptual family resemblances that we may point to. It is doubtful that even family resemblances will work for aesthetic categories such as ‘beauty’ (Sibley 1959).

The worry for aesthetic perception is then that if there are no common perceptually distinguishing features, it remains mysterious how we could perceive beauty or other aesthetic properties. The story one might tell in the non-aesthetic case is not available here, and so a robust account of aesthetic perception should explain why it is that we find such a motley crew of objects beautiful, and what it is that unites all these things in absence of shared perceptible properties (beyond the fact that they are all beautiful). Call this the *diversity* desideratum.

4.2.4 Non-perceptual art forms and objects of aesthetic attribution

Apparently non-perceptual art forms such as literature and perhaps conceptual art pose a problem for aesthetic perception, insofar as we can attribute aesthetic properties to these objects. We may say that a book is beautifully plotted, or that its plot is uplifting, where plots are not plausibly perceptible objects but rather perhaps abstract objects. With respect to conceptual art, we may appreciate the ingenuity or the shrewdness of the idea behind the art (Goldie and Schellekens 2009). If we can attribute aesthetic properties to ideas, then again we may have reason to reject aesthetic perception. Likewise, other objects outside the artworld such as math proofs, and scientific theories may all be said to possess aesthetic properties such as beauty. If these are not perceptual, then the aesthetic properties in this case also cannot be perceptual.

A theory of aesthetic perception should therefore be able to say something to illuminate our practice of attributing aesthetic properties to non-perceptible objects (see Lopes 2014a ch. 9 for discussion of this issue). Call this the *abstracta* desideratum.

4.2.5 The evaluative nature of aesthetic properties

Aesthetic properties, unlike natural kind properties or low-level properties such as colour and shape, are usually taken to be evaluative properties (Sibley 2001c; Levinson 2001; though see Zangwill 2001). Paradigm instances – such as beauty and ugliness – do not serve to merely describe an object but also provide a (dis)favourable evaluation of it. While normativity and felt evaluation may come apart (though see Fingerhut and Prinz 2018), nevertheless the evaluative nature of aesthetic properties is often understood in terms of a felt, affective component.

While there are several ways one might press this disanalogy into an objection to aesthetic perception, one way is to hold that perceptual experience is transparent and aesthetic experience is not, so aesthetic properties are not plausibly perceived (Todd 2014). Perceptual experience is sometimes argued to be transparent in that what it is like to have a perceptual experience can be fully captured by pointing to what the experience is an experience of (Tye 2002). Affective experience, in contrast, is opaque: while it may (or may not) represent an object as possessing certain properties – such as a dog as being dangerous – there is a felt quality to the experience itself (Aydede and Fulkerson 2014; Aydede 2019). Aesthetic properties seem to fall into this latter camp, as there is plausibly something it is like to experience beauty that goes beyond merely attributing the property of beauty to an object. This is arguably part of the hold beauty has historically had on us – it does not just present itself, it captivates and moves us.

In response, one might take a middle ground position and suggest that aesthetic properties are represented in emotional – rather than perceptual – experience (Tappolet 2016). Or, one might argue against perceptual experience being transparent. For example, if we suppose that there are independent compelling reasons to grant that aesthetic properties can be perceived, then they stand as a counterexample to the transparency of perceptual experience. The evaluative nature of aesthetic properties does not count definitively against aesthetic perception. Nevertheless, a view of aesthetic perception should illuminate how it is that aesthetic properties can be both evaluative and perceived. Call this the *affective* desideratum.

4.2.6 Culturally diverse and changing aesthetic ideals

The ability to make aesthetic judgments seems to be universal – ascriptions of beauty and ugliness are common in the folk tales of many cultures from different time periods, most religious ceremonies involve some aesthetic element, and aesthetic adornments feature on many historical objects such as pottery and knives. However, against this backdrop of universality, our aesthetic ideals and judgments often change over time. This is especially true when we look to the realm of human beauty – the actors and celebrities of yesteryears who were considered knockouts at the time would often fall short by today's standards. Similarly, aesthetic ideals and judgments differ from culture to culture. A person or artwork widely judged to be beautiful according to one culture may not receive the same evaluation in another culture.

These shifts in aesthetic judgments put pressure on aesthetic perception, along with any realist view of aesthetic properties. In general, such shifts may suggest that aesthetic judgments are relative in a way that compromises realism about aesthetic properties – the view that

aesthetic judgments are true if and only if the object in question actually possesses the attributed aesthetic properties (McDowell 1983; Pettit 1983; J. O. Young 1997). It seems rather implausible to suppose that some of these judgments (past vs. present, one culture vs. another) are systematically in error, and so such shifts lend plausibility to the view that aesthetic judgments are relative to the culture in which they are made. The *cultural* desideratum thus calls upon an account of aesthetic perception to explain how perception of aesthetic properties is compatible with shifts in aesthetic judgments over time and through cultures.

4.2.7 Aesthetic training

Aesthetic expertise seems to require a great deal of training – we are not born wine connoisseurs or experts in female Renaissance painters. Instead, we acquire these abilities over time, through learning. Training does not sit well with the perceptual view for two reasons. First, it invites a disanalogy with our ability to perceive uncontroversially perceptual properties such as size, colour and shape. We do not seem to need training to perceive these properties, so if training is involved in the case of aesthetic properties then this may count against a perceptual view. Second, if we suppose that such training essentially involves learning about the artist, historical milieu, the production techniques used and so on, then this renders the perceptual view less plausible. It suggests that we draw upon this knowledge somehow to make accurate aesthetic judgments, even though we may not be able to articulate general principles of inference.

The problem of training has led some who appear to endorse aesthetic perception to qualify their views, suggesting that they are using the term ‘perception’ in a broader sense than that denoted by sensory perception. For example, in qualifying his ‘perceptual’ view, Sibley

writes that “*broadly* speaking, aesthetics deals with a *kind* of perception” (2001a, 34 emphasis added).

Sibley contrasts sensory perception with aesthetic perception, noting that while there are similarities between the exercise of taste and the use of the five senses, “there are great dissimilarities too” (2001b, 14). He holds that those with regularly functioning perceptual senses may fail to discern aesthetic properties and that developed aesthetic sensibility is rare among the general population, though this can be cultivated through learning.⁴¹ It is only with aesthetic perception that one must exercise one’s “taste, perceptiveness, or sensitivity, of aesthetic discrimination or appreciation” (2001b, 1).⁴² In the terms of the content debate laid out in the introduction, we might ascribe a view to Sibley whereby aesthetic properties form part of the contents of our overall experience, but not our perceptual experience.

Sir Francis Hutcheson, the eighteenth century ‘inner sense’ theorist, seems to have been similarly motivated to distance aesthetic ‘perception’ from ordinary perception. On the subject of training, Hutcheson writes that:

The proper Occasions of Perception by the external Senses, occur to us as soon as we come into the World; whence perhaps we easily look upon these Senses to be natural: but the Objects of the superior Senses of Beauty and Virtue generally do not. It is probably

⁴¹ This should be considered as a subsidiary argument of its own against aesthetic perception, apart from training, but I will not discuss it at length because it begs the question – if we perceive aesthetic properties then our perceptual faculties are not in perfect working order when we do not perceive these properties, but the argument begins by assuming that our faculties are in perfect working order and so rules out the possibility of aesthetic perception from the beginning. Hutcheson (2008) is also in this territory when he writes that, “I should rather chuse to call our Power of perceiving these Ideas, an Internal Sense, were it only for the Convenience of distinguishing them from other Sensations of Seeing and Hearing, which men may have without Perception of Beauty and Harmony. It is plain from Experience, that many Men have in the common meaning, the Senses of Seeing and Hearing perfect enough [...] And yet perhaps they shall find no Pleasure in Musical Compositions, in Painting, Architecture, natural Landskip; or but a very weak one in comparison of what others enjoy from the same Objects. This greater Capacity of receiving such pleasant Ideas we commonly call a fine Genius or Taste” (p.23, section X).

⁴² There is an alternate interpretation of Sibley, suggested by Dom Lopes (in conversation) that he is endorsing aesthetic perception as learned – genuine – perception. If this is the case, then the view I develop below is complementary to Sibley’s.

some little time before Children reflect, or at least let us know that they reflect upon Proportion and Similitude; upon Affections, Characters, Tempers; or come to know the external Actions which are Evidences of them. Hence we imagine that their Sense of Beauty, and their moral Sentiments of Actions, must be entirely owing to Instruction, and Education.

(Hutcheson 2008, 10)

Again, the role of explicit instruction or education seems to put aesthetic perception in doubt. In order to avoid a retreat to this more broadly experiential view of aesthetic properties, any account of aesthetic perception must provide a convincing story about how it is that training and the perceptual experience of aesthetic properties are compatible. Call this the *training* desideratum.

4.2.8 Lack of aesthetic training

If we focus on the fine arts, then the necessity of some sort of formal training for making accurate aesthetic judgments seems plausible. However, if we focus on everyday aesthetics, or perhaps environmental aesthetics, then formal instruction instead looks implausible.

Everyday aesthetics emphasizes the pervasive nature of the aesthetic: aesthetic properties seem to be part of the fabric of our everyday lives, from choosing what clothing to put on in the morning, to which mug to drink out of. While our everyday experiences are not as intense and focused as our experiences with symphonies and artworks in museums, they are nevertheless

important sources of value (Irvin 2008a; Saito 2001; 2007). Such experiences do not seem to require training.

Appreciation of the natural environment is another potential case. We seem to require no training in order to appreciate the beauty of sunsets or famous actors, or the cuteness of small animals and children, or the ugliness of naked mole rats. This intuition is bolstered by empirical studies showing that babies just a few months old prefer to look longer at more beautiful faces (Samuels et al. 1994; Ramsey et al. 2004; G. Rhodes et al. 2002).

This lack of required training seems to favour aesthetic perception, but an inferentialist may reply that our ability to infer the presence of some aesthetic properties is innate. This view might be less plausible for other reasons – such as attributing inferences to babies – but given the ample evidence of sophisticated cognition in babies the view cannot be dismissed out of hand (R. Wu et al. 2011; Barry, Graf Estes, and Rivera 2015). So the evidence here is not decisive.

However, the contrast between the need for training and the lack of training raises the question of how to reconcile the two: an account of aesthetic perception should be able to explain how aesthetic properties can be both learned and unlearned. This is a sort of addendum to the training desideratum – any account of aesthetic training that is compatible with aesthetic perception should also be compatible with the fact that some aesthetic perception does not involve training. That is, we should prefer a unified explanation of aesthetic perception whereby aesthetic properties can be represented in perceptual experience with and without training. Call this the *un-trained* desideratum.

In summary, we are left without decisive reasons to adopt a genuinely perceptual view of aesthetic properties, and some reason to reject a perceptual view. For those who would defend aesthetic perception, the challenge is then to meet the desiderata specified above.

4.3 Arguments in favour of aesthetic perception

4.3.1 Phenomenal contrast argument

One way of arguing in favour of aesthetic perception is to appeal to the argument from phenomenal contrast (Siegel 2006a; 2010). The argument rests on the intuition that there is a difference in phenomenology between novice and expert, or before and after training. For example, take someone who is about to learn to speak Farsi. It is plausible that after she has achieved some level of proficiency the phenomenology of hearing spoken Farsi is different from her experience prior to learning. Or, take a novice and an expert birder out in the field together. Upon spotting a bird, the expert is able to recognize it as a northern flicker. The novice is not (yet) able to do so. In this case, we may contrast what it is like to see the bird for each, and posit that again, there is a difference in phenomenology.

Granting this intuition, Siegel labels the overall experience of a subject S prior to training ‘the contrasting experience’ and the overall experience of S after training ‘the target experience.’ E1 is the visual experience had by S prior to training, and E2 is the visual experience had by S after training. E1 will thus be part of the contrasting experience, and E2 part of the target experience. Her argument then proceeds along the following lines (Siegel 2010, 101):

- (1) The target experience differs in its phenomenology from the contrasting experience.
- (2) If the target experience differs in its phenomenology from the contrasting experience, then there is a phenomenological difference between E1 and E2.
- (3) If there is a phenomenological difference between E1 and E2, then E1 and E2 differ in content.

(4) If there is a difference in content between E1 and E2, it is a difference with respect to [high-level] properties represented in E1 and E2.

(5) Therefore, there is a difference with respect to high-level properties represented in E1 and E2.

(Siegel 2010, 101)

While Siegel's original argument does not focus on aesthetic properties, it is equally applicable to such properties. For example, (Stokes 2014, 13–14; 2018a) contrasts one's overall experience of a painting before one can recognize it as an impressionist painting and after one can recognize it as such, suggesting this contrast might be used to argue that we perceive the high-level property of 'being an impressionist painting.' He also contrasts one's experience of a dancer's movements before and after learning to recognize gracefulness, suggesting this points to a view where we represent the high-level aesthetic property 'gracefulness' in perceptual experience.

The application of the method of phenomenal contrast to aesthetic properties has received some criticism (Nanay 2015 section 4.5; Logue 2018). Logue's (2018) first criticism of the argument is on grounds similar to criticisms of non-aesthetic applications of the argument (Logue 2013; Price 2009; Connolly 2014). If we grant that attentional differences can result in phenomenological differences while holding the contents of perception fixed, then we need not posit a difference in content to explain the contrast: it may be explained by holding that experts learn to allocate their attention to different aspects of their perceptual experience than novices do. There is some empirical evidence for this in the aesthetic realm where, in viewing paintings,

experts exhibit systematically different patterns of visual attention from those of novices (Francuz et al. 2018).

Logue's second criticism is specific to the aesthetic application of the phenomenal contrast argument. She argues that the difference between the novice and the expert in recognizing a ballet dancer's pirouettes as graceful may be best explained instead by a phenomenological difference in emotional states. She points out that our aesthetic judgments are reliably connected with broad emotional responses. In the case of gracefulness, this might take the form of experiencing pleasure. Her conclusion is that, in the face of these plausible alternatives, the aesthetic application of the phenomenal contrast argument is inconclusive in establishing aesthetic perception.⁴³

4.3.2 Can we appeal to empirical data?

While there is a growing field of empirical aesthetics, it has not provided us so far with a means of adjudicating the debate over whether aesthetic properties can be perceived. Some have raised doubts over the prospects of appealing to empirical data to answer this question. For example, Nanay (2015 section 4.5) points to studies of patients who suffer from unilateral neglect. Such patients are able to consciously experience an object and recognize what the object is used for, but they do not consciously experience the object's low-level properties such as its shape and size. This sort of case, he argues, is helpful for determining the contents of perception in non-aesthetic cases, but it is unlikely that we will be able to find patients who continue to attribute aesthetic properties but lack an experience of the object's low-level properties. At least,

⁴³ The variant on the phenomenal contrast argument proposed by (Stokes 2018a) to argue in favour of aesthetic perception can also be criticized along these lines. While one may grant there is a change in overall phenomenology, this may be via the emotions rather than perception.

there are no known such cases, and so this same strategy cannot be used to argue in favour of the perception of aesthetic properties.

One might wonder whether the common mechanism argument I discussed in chapter two and employed in chapter three to argue for the representation of some high-level properties in perceptual experience might be of use here. In principle, it could be of use: there is no obstacle to setting up experiments to gather data on how fast we can detect aesthetic properties, or whether they exhibit adaptation or pop-out effects. If we had such data then this would serve to inform the argument for or against aesthetic perception in the way discussed in chapter one.

However, such data is sadly lacking in most cases, though there is some work with regards to facial adaptation and judgments of beauty. Rhodes et al. (2003) found that after subjects were shown distorted adaptor faces (e.g. eyes very close together) they shifted their judgments of the attractiveness of undistorted faces towards the direction of the distortion (e.g. judging faces with the eyes slightly close together as more attractive). While this study is useful in supporting the point that judgments of attractiveness track prototypicality (see section 4.5), it is not quite the sort of experiment needed to show that aesthetic properties exhibit adaptation effects. Such an experiment would need to show that after adaptation to very beautiful faces, subsequent ‘normal’ faces are rated more ugly than they would be rated without exposure to the adaptor face. The common mechanism argument is not (yet) useful in arguing for aesthetic perception.

4.3.3 How to move forward: providing a positive account

One way of providing an argument in favour of aesthetic perception is to develop an empirically grounded account of *how* aesthetic perception occurs. This does not serve to refute other views, but instead raises the bar for defenders of these other accounts. Inferentialists and those who hold that aesthetic properties are only represented in the non-perceptual components of experience must either explain why the account does not really qualify as perceptual, or find a means of discounting the empirical data.

In section 4.5 below I provide just such an account, which serves as an argument in favour of aesthetic perception. First, however, I review another prominent attempt to explain how aesthetic perception occurs: cognitive permeation.

4.4 Aesthetic perception by cognitive permeation

One proposal for how we come to perceive aesthetic properties is that our background knowledge permeates our perceptual experience. Cognitive permeation (also commonly known as cognitive penetration) is hypothesized to occur when the contents of perceptual experience are altered in some way by one's cognitive states, such as beliefs and desires (Pylyshyn 1999; Macpherson 2012; Stokes 2013; 2018b; Briscoe 2015; Raftopoulos 2009; Newen and Vetter 2017; Gross 2017; Siegel 2012; 2013).⁴⁴

While there is no uncontroversial way of defining the phenomenon, the alteration of perceptual experience by cognition is usually said to take place via an internal causal connection, and the relationship between the permeating state and the resulting perceptual experience is

⁴⁴ Though traditionally the term is 'cognitive penetration,' I follow Becco Copenhaver (personal conversation) in using the term 'cognitive permeation' as it paints a less gendered picture of the relationship between perception and cognition.

sometimes said to be quasi-logical or at least semantically coherent (Stokes 2013; Zeimbekis and Raftopoulos 2015b, 27–32). This rules out cognitive effects on perceptual experience that are the result of directing our attention to some things over others: while I might attend to the chair where I believe my friend to be sitting, and indeed come to perceive her sitting there, this does not count as cognitive permeation because the mediating role of attention fails to count as an internal connection (though see Mole 2015; W. Wu 2017; Stokes 2018b). And if I should instead discover that a stranger is sitting in the chair, though my cognition has altered my perceptual experience it has done so in a way that is semantically arbitrary – the content of my belief does not bear a coherent relation to the resulting content of my perceptual experience. Put a slightly different way, we may say that perceptual experience is not counterfactually sensitive to the content of the belief (Siegel 2012). If I had believed there was a crocodile sitting in the chair, I would still perceive the stranger to be sitting there.

Cognitive permeation has been suggested to be a mechanism that allows background knowledge and training to play a role in our aesthetic judgments while nevertheless preserving a perceptual view (Lamarque 2010; J. Margolis 2000; 1998; Nanay 2015; Stokes 2014; 2018a; Anna Bergqvist and Cowan 2018). This view has received limited criticism (Danto 2001a; 2001b). Cognitive permeation has also been a popular interpretation of the perceptual ‘contextualist’ view Kendall Walton developed in his (1970) paper, ‘Categories of Art.’ Walton’s view is nevertheless key to understanding how we can come to represent aesthetic properties in perceptual experience.

4.4.1 Walton's categories of art

In 'Categories of Art,' Walton first defends the psychological thesis that what category we perceive an artwork in alters what aesthetic properties we take the work to have. On Walton's view, to perceive a work in a category of art is to perceive certain non-aesthetic perceptible properties (hereafter 'non-aesthetic properties') as standard, where these standard properties are perceived as somehow unified or connected into an overall 'gestalt.' A standard property is one that most works in that category possess. For example, the standard colour palette and brushstroke style of impressionist painters is unified in such a way that an observer may perceive the work as an impressionist painting. In making such categorizations, contrastandard properties are also relevant. A contrastandard property is one that the work possesses that tends to disqualify it from membership in that category, though it may not on its own be sufficient to do so. Perception of non-aesthetic properties as contrastandard then indirectly contributes to categorizing a work. For example, a very dark colour palette may count against a work's qualifying as an impressionist painting. Finally, all other non-aesthetic properties of a work are perceived as variable, where a variable property neither counts against or for a work's inclusion in a category.

Perceiving the non-aesthetic properties as standard, contra-standard or variable in turn affects what aesthetic properties the work appears to have. Walton argues for this point by way of examples. Given that sculptures are standardly static, a sculpture with a twitching element (a contrastandard property) would be perceived as shocking. Standard properties may contribute to a sense of order or stability in the work, such as in the case of the first movements of sonatas in classical music, which standardly have an exposition-development-recapitulation structure. His most widely-discussed example is that of Picasso's painting *Guernica*, which may be said to

possess aesthetic properties such as ‘violent’ and ‘dynamic’ when viewed as a painting. But imagine that there is a society that does not have a tradition of painting but rather an artform labelled ‘guernicas’ that is composed exclusively of artworks that are identical to Picasso’s *Guernica* in colour and line but vary topographically, with sections of the artwork rising and falling in various patterns of bas-relief. Upon viewing Picasso’s *Guernica* for the first time, members of this society would mistakenly categorize it as a guernica and likely judge it to be ‘dull’ and ‘lifeless’ given that it is perfectly flat. Walton explains this difference in judgments in terms of which properties are standard to the respective categories. While flatness is a standard property with respect to the category of painting, it is contrastandard with respect to the category of guernicas. And while colour and line are variable for paintings, they are standard for guernicas.

The psychological thesis only gets Walton so far however, as it does not give him the resources to say when someone is mistaken in their aesthetic judgment. If one person perceptually categorizes an artwork as a guernica, and the other as a painting, then for all Walton has said both may have accurate aesthetic experiences. For this reason, Walton goes on to defend the normative thesis that what category or categories a work actually belongs to depends in part on facts about the work’s production history, the artist’s intentions, and the society in which it was produced.

4.4.2 ‘Categories of Art’ as cognitive permeation

Because of Walton’s emphasis on the importance of these extra-perceptual facts, several subsequent contextualists have interpreted his view as one that that invokes cognitive

permeation, or has been adopted by those who uphold a cognitive permeation view. For example, Lamarque writes that “[a]ll perception is informed by background knowledge. [...] What Walton’s argument establishes so powerfully is that our aesthetic responses are thoroughly determined by our beliefs about what kind of thing we are looking at” (Lamarque 2010, 132). For other similar interpretations see (Hopkins 2005; Todd 2014).

Stokes (2014) argues that cognitive permeation can explain an underdeveloped aspect of Walton’s view, and so lends greater plausibility to the claim that we perceive aesthetic properties. This is what Stokes labels the ‘expertise-to-perception effect’ – how it is that aesthetic expertise can come to change the contents of perceptual experience.

Walton does not provide a detailed picture of how expertise leads such to a perceptual change, stating only that:

Perceiving a work in a certain category or set of categories is a skill that must be acquired by training, and exposure to a great many other works of the category or categories in question is ordinarily, I believe, an essential part of this training. (But an effort of will may facilitate the training, and once the skill is acquired one may be able to decide at will whether or not to perceive it in that or those categories.)

(Walton 1970, 366)

Stokes proposes to understand this training, and so how expertise can come to influence perception, in terms of cognitive permeation. Aesthetic expertise will involve acquiring knowledge that helps people place artworks into the correct categories of art, such as ‘impressionist’ or ‘in the style of VanGogh.’ This knowledge of the correct category either

cognitively permeates our perceptual experience in the moment or over time, so that we represent the work in perceptual experience as a member of that category of art.⁴⁵ The perceptual representation of a work in a given category of art will then affect what aesthetic properties we perceive the work to have, just as Walton proposed.

4.4.3 Against ‘Categories of Art’ as cognitive permeation

This proposal has the virtue of being able to meet the training desiderata. It explains how training is compatible with aesthetic perception: training provides us with the background beliefs regarding categories of artworks that then permeate perceptual experience and allow us to represent these properties in perceptual experience.⁴⁶ Nevertheless, it does not fare as well when it comes to explaining how we come to represent aesthetic properties themselves in perceptual experience, what Stokes terms the perception-to-aesthetic reaction effect. We can come to represent categories of art such as ‘impressionism’ in perceptual experience via background knowledge, but it remains to be explained how this results in a change in the aesthetic properties we are able to represent in perceptual experience.

Here one potential move is to suggest that, just as we can form perceptual categories of art through training, so too can we do the same for aesthetic properties. Stokes writes (though it is not clear if he ultimately endorses this view): “when one learns what a telephone is, one may perceptually represent telephones. Analogously, if one learns what gracefulness is – how to

⁴⁵ Stokes also considers a more conservative view in which we do not represent categories of art in perceptual experience, but where such knowledge permeates perceptual experience and causes the perception of aesthetic properties.

⁴⁶ However, I have argued elsewhere (Ransom 2020) that the cognitive permeation interpretation goes against Walton’s criteria for what counts as a perceptually distinguishable category. On Walton’s original formulation, background knowledge is neither necessary nor sufficient for perceptually categorizing artworks, and the categorization must be made by means of perceptually distinguishable features alone.

recognize being graceful as such – then one can perceptually represent, for example, a ballet dancer as graceful” (Stokes 2014, 12).

But this proposal does not work. Aesthetic properties are not plausibly like artifactual kind properties such as ‘telephones.’ As discussed in section 4.2.3, the objects that possess a common aesthetic property may vary widely in their perceptual features: the graceful neck of a swan, the graceful pirouette of a dancer and a cellist’s graceful execution of a solo have little in common perceptually other than all possessing the aesthetic property of gracefulness. Learning to recognize gracefulness does not seem at all to be like learning to recognize telephones, given this wide variation in non-aesthetic properties.

If one instead answers the question of how we come to perceive aesthetic properties by positing that acquired knowledge of aesthetic properties directly penetrates experience, then one needs to then explain how we can come to have the knowledge that these aesthetic properties are instantiated in the first place. Such knowledge is either acquired by inference, testimony, or perception.⁴⁷

If one claims that this original knowledge is acquired via non-perceptual means via the learning and application of (perhaps unconsciously held) principles or rules of inference, then the challenge of aesthetic particularism must be answered. There is a conspicuous lack of any general rules for inferring aesthetic properties from non-aesthetic properties, and this is not for lack of trying (Sibley 1959; Bender 1995; Strahovnik 2004; Lopes 2014b). For a particular painting it might be true that its having a bright yellow spot makes it dynamic, but a bright yellow spot on another painting might have the effect of making it unbalanced, clumsy, or garish.

⁴⁷ There are more options here, such as a priori intuition. I leave it to others to explore whether these options are plausible.

So choosing this route requires finding and articulating plausible general principles of inference, which is perhaps a Sisyphean task.

Appealing to inductive inference or testimony (Cavedon-Taylor 2017) also does not help. Our knowledge that certain aesthetic properties are instantiated is likely sometimes derived from these sources. Based on our knowledge of a band's past repertoire, we can form a reasoned belief that their next song is likely to be melancholic. But then we may ask how we arrived at our past judgments, and the answer must be inference or perception. The same issue arises for testimony. A trusted expert tells us that a painting is dynamic and jarring. However, the expert must herself have formed her judgment either on the basis of inference or perception.

If one claims instead that the knowledge is acquired via perceptual means, then this renders the inclusion of cognitive permeation at best redundant, at worst incoherent. On this view we perceive aesthetic properties in order to form beliefs that the aesthetic properties are instantiated, which then permeate our perceptual experience. Perhaps this view can be extended by positing that some aesthetic properties require no training or background knowledge to be perceived in at least some objects. Only more complex categories, or perhaps only a subset of aesthetic properties, require cognitive permeation via background knowledge. If this option is chosen, then cognitive permeation still ultimately requires a defence of aesthetic perception without cognitive permeation.

The cognitive permeation view is then incomplete in two ways. First because it does not explain how we are able to come to perceive aesthetic properties in cases where aesthetic perception seems not to require training or background knowledge. In such cases, cognitive permeation is implausible, and so a separate explanation of aesthetic perception is still needed. Second, the view is incomplete because even in those cases where aesthetic perception occurs

via cognitive permeation, the perception-to-aesthetic reaction effect has not been explained.⁴⁸

While it provides a clear account of how expertise can allow us to perceive artworks as belonging in a given category of art, it does not explain how this allows us to perceive aesthetic properties themselves.

A way of addressing this second issue is to deny that any further explanation of the expertise to perception effect can be given. On one way of interpreting Walton's (1970) proposal, aesthetic properties may be said to supervene on an object's perceptible non-aesthetic properties, where some of these properties are perceived as standard, contrastandard and variable properties (Stokes 2014, 8 fn. 14). If this is the case, then perhaps nothing more is needed to explain how we come to perceive aesthetic properties, because there is nothing more to say – we should take their supervenience to be explanatorily basic. Once we allow that we perceive the artwork in a category then we get aesthetic perception for free. The perception-to-aesthetic reaction effect is then wholly explained by the acquisition of background knowledge of categories of art, which are the only beliefs said to permeate perceptual experience.

However, I think that we can and ought to do better in terms of offering a complete account of the perception-to-aesthetic reaction effect. In what follows I provide at least the beginnings of an explanation for how it is that perceiving certain non-aesthetic properties as standard, contrastandard or variable thereby influences or determines what aesthetic properties we perceive the work to have. I also provide a fully developed view – which stands as a competitor to the cognitive permeation account – of how it is that our perception of which properties are standard, contrastandard or variable is influenced by training and expertise. In

⁴⁸ It is also worth noting that there has been significant recent pushback on the empirical studies that purport to demonstrate that cognitive penetration occurs, criticizing the methodology of these studies and thus casting doubt on whether there is adequate evidence for its occurrence (Firestone and Scholl 2014; 2016; Machery 2015).

keeping with the account developed in chapter three, perceptual learning explains how we can come to form perceptual categories for certain categories of art and so come to represent artworks as members of a given art category in perceptual experience. By understanding the structure of perceptual categories, we can begin to understand how we perceive aesthetic properties, and so fill in the perception-to-aesthetic reaction effect. Such properties are not themselves identified via perceptual categories. Rather they are identified as the result of the fluency with which an object is categorized in perceptual processing.

While this proposal is consistent with Walton's account, and draws on several aspects of the framework laid out in Walton (1970) it also goes beyond it in several respects, and there are elements of this account that the present-day Walton (2020) disavows. Therefore, I call this view 'Waltonian' perceptualism.

4.5 Waltonian perceptualism

Waltonian perceptualism contains three main strands that will be elaborated on below. First, perceptual learning explains how it is we are able to categorize objects, including artworks, in perception. Positing that such learned perceptual categories exhibit a prototypical structure sets the stage for understanding how aesthetic perception emerges, as it provides the key link to the second main strand of the account: the fluency of perceptual processing. Recent empirical work links processing fluency to our perceptions of beauty and ugliness, and so provides the beginning of an answer to the perception-to-aesthetic reaction effect. While aesthetic properties are not themselves perceptual categories, they are attributed to objects via the affect produced by perceptual processing. Finally, I put forth an account of perceptual expertise whereby it can be

understood in terms of possessing perceptual categories that accurately reflect the world. This gives us the resources to say when aesthetic perceptions (and concomitant judgments) are mistaken or veridical. The picture that emerges is one where we can truly perceive beauty (and perhaps other aesthetic properties), but only if we are perceptual experts in a given category. For those who are not experts, apparent instantiations of beauty are a perceptual illusion.

4.5.1 Perceptual learning

While it was long thought that the perceptual system remained relatively fixed after a short ‘critical period’ in early development (Wiesel and Hubel 1963; Hubel and Wiesel 1970), there is a wide body of evidence that attests to its plasticity well into adulthood (Hooks and Chen 2007). One branch of study of this plasticity is perceptual learning. Following Gibson (1963) and Goldstone (1998), perceptual learning involves structural and functional changes in the perceptual system, due to repeated exposure to a stimulus, that result in a change in perceptual experience (see also Connolly 2019; Kellman and Garrigan 2009). Importantly, perceptual learning is perceptual: the changes brought about must make a difference to how something is perceived, or whether it is perceived at all. The changes must be brought about as a result of learning, which discounts changes due to lesions and aging, for example. Learning must also take place as a result of repeated exposure to or practice with the target stimuli.

Some perceptual learning can be wholly explained in terms of low-level properties. For example, the ability to discriminate the direction of motion of dots on a screen improves with extensive perceptual practice with the task (Ball and Sekuler 1987; 1982; Matthews et al. 1999). In such cases, the improvement is uncontroversially perceptual – experimental design ensures

that subjects are not getting better just by guessing or inferring direction of motion. Nevertheless, in such cases we need not posit that improvements in subjects' abilities are explained by the representation of any high-level properties in perceptual experience.

However, there is also empirical work that suggests that we can come to represent high-level properties in perceptual experience as the result of perceptual learning. The same processes involved in the perceptual learning of low-level properties are hypothesized to allow for the formation of perceptual categories. Here I will discuss two such processes: attentional weighting and topographical imprinting.

Attentional weighting is the process whereby the features of a stimulus come to be weighted more or less heavily in attention, depending on their relevance for performing tasks, including the task of categorizing objects (Tanaka and Taylor 1991; Goldstone and Steyvers 2001). Experts and novices with respects to object categories do exhibit systematically different attentional patterns, such as radiologists with expertise at detecting tumours via x-rays (Drew et al. 2013; van der Gijp et al. 2017), and artists with expertise in assessing the composition and other aesthetic features of paintings (Vogt and Magnussen 2007).

Such learned attentional weighting becomes relatively stable for categorization when a certain level of proficiency is achieved (Richler, Wong, and Gauthier 2011). One source of evidence for this stability comes from studies in which experts (but not novices) with a given category are unable to ignore features that are diagnostic for categorization but irrelevant for a transitory task (Gauthier et al. 2003; A. W. Young, Hellawell, and Hay 2013). Experts at perceptually categorizing a given object cannot reweight the features of an object at will – at least in the short term.

This learned attentional weighting in turn influences which features of the stimulus are most salient (or present at all) in our perceptual experience. Some evidence for this comes from a phenomenon known as ‘categorical perception,’ whereby it is more difficult to perceptually distinguish between objects belonging to the same learned category than it is to distinguish between objects of different learned categories, controlling for actual physical differences between objects (for a review see Goldstone and Hendrickson 2010). The best known categorical perception effect is with speech phonemes, where – despite their being exposed to phonemes altered along a physically smooth continuum from /be/ to /ge/ – subjects’ performance at a discrimination task and their subjective experience tracked abrupt and ‘stepwise’ sound transformations directly from /be/ to /ge/ (Liberman et al. 1957). It is partly for this reason that attentional weighting cannot be dismissed as an extra-perceptual – the differential weighting of object features for categorization results in a relatively stable and enduring change in perceptual experience.

Topographical imprinting is a process whereby the detection of isolated features becomes unitized so that they are detected as a single unit, and more abstract relations – sometimes called ‘second-order relations’ – between these features are encoded.⁴⁹ Topographical imprinting is somewhat like creating a map of the object to be detected. Maps also abstract away from details to depict only those features relevant to navigation, and they also represent the spatial relationships between these features. An example of topological imprinting is face detection. Face processing is thought to be at least in part concerned with the perception of second-order relations between facial features. Not only do we detect noses, mouths and eyes in specific

⁴⁹ In fact, unitization is a separate process hypothesized to occur in perceptual learning. The main difference between topographical imprinting and unitization is that in the former case the features might have already been grouped together to some extent, whereas in the latter, they are separate prior to learning.

patterns (the eyes above and perpendicular to the nose above the mouth) but we are also sensitive to the distance between these features, and use these relations to recognize individuals.

One source of evidence for this phenomenon comes from research on perceptual expertise, where experts with respect to a given domain such as faces, cars, birds, or special lab-manufactured objects cannot help processing these objects holistically. Experts (but not novices) have difficulty performing tasks that require perceptually isolating individual attributes, or that distort or obscure the second order relations between the individual features (Gauthier, Tarr, and Bub 2010).

Together, these two processes of perceptual learning can result in the creation of perceptual categories that correspond to high-level properties. When objects are categorized in perception, we can thereby come to represent them as possessing the corresponding high-level property.

On some accounts, the resulting perceptual category possesses a prototypical structure. Prototypes encode in some way the central tendencies of category members, and use the encoded information in order to categorize objects (Rosch et al. 1976; Rosch 1973; Mervis and Rosch 1981; Rosch and Mervis 1975). While the exact details of what is encoded varies on different accounts of prototypes (E. E. Smith and Medin 1981), in general they are perceptual features of objects that are typical of the category. For example, dogs have tails, four legs, fur, elongated snouts, and make barking noises, amongst other characteristics. Prototype theory, in contrast to the classical theory of concepts, does not hold that there is a set of necessary and sufficient conditions for category membership. Instead, the encoded set of central tendencies is used to determine membership, where this involves possessing a sufficient number of the relevant features to clear some pre-determined threshold. While Boston terriers have short snouts, and

Basenjis do not bark, they still possess enough other typical features to be categorized as dogs. In many prototype models, the typical features are weighted according to diagnosticity: while the length of an animal's snout predicts to some degree whether it is a dog, making a barking noise is a much better diagnostic, and so should be 'weighted' more heavily in categorization (E. E. Smith and Medin 1981).

Given this description of prototypes, it is easy to see how one might make the connection to perceptual learning: the process of attentional weighting of features is plausibly understood in terms of their relative diagnosticity for category membership, and the grouping of these features together in topographical imprinting or unitization can be understood in terms of the formation of the central tendencies of a category.

Putting perceptual learning and prototype theory together provides us with an alternate explanation of how it is we can come to form perceptually distinguishable categories of art, one that is compatible with Walton's initial characterization, but which fills in the missing empirical details. On Walton's account, to perceive an artwork in a category is to perceive certain of its non-aesthetic properties as standard, where these properties are unified or grouped into an overall gestalt. The perceptual learning process I have described above can be understood as a way of adopting and expanding on this claim: those features weighted in terms of diagnosticity for category membership will be standard features, and their unification into a gestalt can be understood via topographical imprinting.

The account goes beyond Walton's in several ways. First, it has built into it the notion of graded standard properties: some standard properties will count more or less for category membership, depending on their attentional weights. Second, it allows for the learning of the stimulus structure, and for these second order relations between features to themselves be

relevant to categorization. Not only can we perceive non-aesthetic properties as standard, but also the characteristic relations between them. Third, it allows us to understand why it is that aesthetic training so centrally and necessarily involves exposure to multiple works in a given category. Without such exposure, we cannot develop a perceptual category of art; perceptual learning requires exposure to many exemplars, and perceptual expertise requires an unbiased sample of such exemplars (see section 4.5.3).

Perceptual learning stands in contrast to cognitive permeation (see chapter five). Knowledge of which category of art a work belongs to is not necessary for the creation of a perceptual category, though it can serve as a catalyst to perceptual learning. As such, it unites our aesthetic appreciation of art objects and our aesthetic appreciation of non-art objects: all involve perceptual learning, but the explicit training that occurs in art education can accelerate this learning.

Learned background knowledge can guide our attention to the properties that ought to be more heavily weighted, or to the properties that ought to be differentiated or unitized. However, the knowledge itself does not accomplish the change in the perceptual prototype directly (this is why it does not count as cognitive penetration – the background knowledge itself does not penetrate perceptual experience). Rather, the change must happen via perceptual learning – the subject must either be repeatedly exposed to the object, or in some simple cases just differently orient her attention a few times. This will depend on the complexity of the object and the discriminations to be made.

While this view does not yet provide an explanation of how we come to perceive aesthetic properties – the perception-to-aesthetic reaction effect – it nevertheless contains the seed of the solution. Prototypicality is an important factor that determines the fluency with which

an object is processed in perception, and recent empirical work links processing fluency to our perceptions of beauty and ugliness. The prototypicality of objects as they pertain to a perceptual category will then influence our perceptions of these aesthetic properties.

4.5.2 Fluency hypothesis

The fluency hypothesis was first proposed by Reber, Schwarz and Winkielman (2004). The hypothesis can be broken down into two main parts: (1) an empirical hypothesis that the fluent processing of a perceptual stimulus gives rise to positive affect, and (2) a metaphysical claim that this positive affect can be identified with aesthetic pleasure or beauty.⁵⁰ In what follows, however, I do not take on board their metaphysical claim. While a full discussion of the metaphysics of aesthetic properties is beyond the scope of this chapter, I take it that an objective account of beauty – such as that offered by Mole (2016), for example – is compatible with the view I set out in this chapter. Positive affect may merely be what allows us to detect, and so represent in perception, the property of beauty.

Metaphysics aside, empirical work on the fluency of perceptual processing holds the answer to understanding how we come to represent at least some aesthetic properties in perceptual experience. Perceptual processing fluency is the ease with which an object can be perceptually processed. It usually gives rise to a felt, affective component – with relative ease producing positive affect, and relative difficulty producing negative affect. There are two main

⁵⁰ Reber et al (2004) define beauty as interchangeable with aesthetic pleasure. However, on my account of beauty offered below, aesthetic pleasure cannot be identified with beauty as there is an external condition on its veridical instantiation as well. Whether something is beautiful or not is objective, on my view, though it relies on the abilities of the subject in order to be detected.

kinds of factors that influence processing fluency.⁵¹ First, some stimuli can be processed more easily than others in virtue of properties of the object that do not require that one be able to perceptually categorize it. These include those that have been singled out by proponents of objective accounts of beauty as responsible for our aesthetic attributions. For example, symmetry has long been pointed to as a factor that contributes to beauty (Arnheim 1954; Birkhoff 1933; Gombrich 1980; Humphrey 1997; Lloyd 2010). Humans generally have a preference for symmetric over non-symmetric patterns, all else equal (Humphrey 1997; Reber and Schwarz 2006). In the case of faces, symmetry (or lack thereof) is one of the factors that influences perceived attractiveness (Gangestad, Thornhill, and Yeo 1994; G. Rhodes et al. 1998; G. Rhodes, Sumich, and Byatt 1999). The fluency explanation is that symmetric stimuli are easier to process because they contain redundant information. In the case of a stimulus that is symmetrical about a vertical axis (the left side and right side are ‘mirror images’) the perceptual system only needs to process the one half to get all the information it needs (Reber, Schwarz, and Winkielman 2004; Reber 2002). In general, stimuli with less information in need of processing – where informational content is understood as a broader notion than whether the object is symmetrical – will lead to greater processing fluency. For an examination of this principle at work in judgments of Cubist paintings, see (Nicki, Lee, and Moss 1981).

The second kind of factor that influences processing fluency has to do with properties of the object that depend on a person’s ability to perceptually categorize the object. These correspond to the discussion of perceptual learning in the previous section. The learning of the characteristic structural relations between features of a stimulus – in perceptual learning what is

⁵¹ Here my analysis diverges somewhat from Reber et al’s: they do not discuss the distinction in terms of categorization, but rather in terms of properties of the object vs properties that depend on a subject’s previous experience with the object.

referred to as topographical imprinting – enhances processing fluency (Buchner 1994). Structure, in this context, can be thought of as the pattern, grammar, or organizing principles for constructing a stimulus. For example, in one experiment subjects who learned musical sequences constructed according to a particular grammar not only preferred the sequences that they had been exposed to during training, but also novel grammatical sequences over novel ungrammatical sequences (Sollberger and Reber 2004). Similar effects have been found for letter strings following (or violating) an artificially constructed grammar and for visual patterns adhering to (or deviating from) a standard format (Gordon and Holyoak 1983; Manza, Zizak, and Reber 1998; B. R. Newell and Bright 2001).

There is also some evidence that the closer the stimulus is to our perceptual prototype for that category of object – the higher the preponderance of standard features relevant to diagnosticity for category membership – the more fluently it will be processed. We have demonstrated aesthetic preferences for prototypical faces, furniture, paintings, colour patches, music, dogs, watches, and birds, amongst other objects (Langlois and Roggman 1990; G. Rhodes and Tremewan 1996; Martindale and Moore 1988; Hekkert and Wieringen 1990; Whitfield and Slatter 1979; Repp 1997; Halberstadt and Rhodes 2003; 2000).

The fluency hypothesis thus provides an explanation of how it is that our perceptual experiences of at least two aesthetic properties – beauty and ugliness – come about: we process objects more or less fluently, which produces positive or negative affect. This affect causes us to perceptually experience the object as having the aesthetic property of beauty or ugliness. The account thus provides us with a way of filling the explanatory gap between perceptual categories and the perceptual experience of some aesthetic properties. The reason why the aesthetic properties of beauty and ugliness may depend in part on which category of art a work is

perceived in is because the category influences the fluency with which the object is perceptually processed. The account also extends to objects that are not artworks: objects that belong in natural kind categories (e.g. 'dog,' 'horse') and artifactual kind categories (e.g. 'toaster,' 'car') can also be processed more or less fluently, and so also be said to possess aesthetic properties.

However, this account does not yet have the resources to determine whether the perceptual experience of these aesthetic properties is veridical or not. We might construct a troublesome case where one person experiences an object as beautiful and the other experiences it as ugly because they each have differing histories of perceptual learning. This might lead one to think that on the fluency hypothesis beauty is not an invariant property of an object, but rather is relative to each subject.

To resolve the threat of anything-goes relativism, one final element must be added to the account. We must find a way to say when or how a person's experience with an object will lead to accurate or apt aesthetic judgment. Note that this is a different and more involved project than Walton's. In formulating his normative thesis that the correct category of art depends on non-perceptible features such as artist intentions and method of production, amongst other socio-historical factors, Walton was looking only for a way to say when our perceptions of artworks as belonging to categories is correct.

By contrast, here there is an additional issue. On the understanding of perceptual categories of art as prototypes, there may be cases where we are able to perceive an artwork in the correct category but nevertheless have an 'illusory' experience of an artwork as possessing a given aesthetic property that it does not in fact possess because our prototype itself is biased in some way. Perhaps someone's perceptual category of 'Picasso paintings' was formed exclusively through exposure to works from his blue period. Upon perceiving a painting from his rose

period, this person may wrongly attribute an aesthetic property to the painting that it does not possess (perhaps viewing it as garish, for example), simply because it is a psychological – though not actual – outlier.

In order to address this issue, I develop a sketch of what it is to be a perceptual expert. My proposed definition of perceptual expertise provides us with the tools to separate veridical perception of aesthetic properties from illusory perceptions.

4.5.3 Perceptual expertise

My aim in this section is not to provide a complete account of expertise. There are many skills that go into making an expert that are non-perceptual (Montero 2016). What I am interested in is providing an account of what it means to be a perceptual expert with respect to some category of object or domain (see also Ransom 2019). My offered definition is as follows: a person is an aesthetic perceptual expert with respect to a given category C if and only if she is able to reliably perceptually categorize objects in C, where her internal representation of C (the perceptual category or prototype of C) suitably approximates the true population average for C.

The requirement that the classification occur via perceptual categorization rather than non-perceptual concept application makes it such that the expertise be truly perceptual – the expert is not merely good at deducing or inferring that something belongs to a given category by relying on background knowledge, she perceives it as such. While she may also possess expertise of this second type, it is not necessary for perceptual expertise.

Reliable perceptual categorization is necessary because to perceptually process an object in the wrong correct category will lead to different fluency effects that will likely result in

erroneous aesthetic ‘perceptions.’ This failure is perhaps a common way in which aesthetic perception can be in error – we may incorrectly categorize an object and so perceive it as a highly atypical member of the incorrect category. The children’s story of ‘the ugly duckling’ nicely illustrates this sort of mistake. All the other ducks perceived the ‘duckling’ to be very ugly, but it turns out that they miscategorized the ‘duckling,’ which was in fact a swanling. As a member of the proper category, the swanling may have been beautiful. The ducks apparently were not perceptual experts with respect to the category swanling (or else they would not have made the mistake in the first place), and so were not in a position to perceive its true aesthetic properties.⁵²

This requirement is also compatible with the claim that artworks can and often do belong to multiple categories – for all I have said here, one may be able to reliably perceptually categorize an artwork in multiple categories, or one may be a perceptual expert in one category to which the work in fact belongs but not with respect to another.⁵³

However, the requirement that such categorization be reliable is a necessary but not sufficient requirement for accurately perceiving beauty (at least, the sort of beauty that is category-relative). The expert must be able to place the object in the right category in order to judge its beauty accurately, but doing so does not yet guarantee that she will judge it accurately, for the reason outlined above; sometimes our perceptual categories themselves may be biased, based on a skewed sample.

⁵² In the canonical version of the story, the ducks realize their mistake when the ‘ugly duckling’ grows up into a beautiful swan. We might understand this by positing that the ducks do have a perceptual category for full-grown swans, or that swans possess enough ‘objective’ beauty-making properties so that even those who have never seen a swan before can find them beautiful.

⁵³ Here I adopt Walton’s normative thesis in order to determine which perceptual category an artwork in fact belongs to.

This is the motivation for the condition that the perceptual category or prototype must accurately mirror the true population average for members of a given category. What this condition amounts to is the requirement that the evidential weights assigned to each feature must (within some reasonable range) accurately reflect how relevant that feature is to diagnosing category membership. Because almost all Dalmatians are white with black spots, colour will be a highly diagnostic feature in this case. Colour is not a strong diagnostic for dogs, however, whose colouring varies over the different breeds, and so colour should not be too heavily weighted in this case. The result is that the evidential weights track the true population statistics for those features, where the weight corresponds to the average prevalence of the feature in a population. This is distinct from the requirement that the prototype track all the statistically relevant features of a population – miniscule spots on the underside of a dog may be highly diagnostic of belonging to a certain breed, but a prototype need not track this feature. Rather, the requirement applies only to those features or relations between features that are in fact used to diagnose membership.⁵⁴

Requiring that diagnostic features be weighted according to their actual prevalence is what allows us to distinguish between veridical and illusory cases of aesthetic perception. We might perceptually experience some object to be beautiful due to its relatively fluent processing, but if we have only been exposed to a skewed sample of exemplars, then the prototype we will have formed will likely be highly biased towards an average that is quite distant from the true population average. In such cases our perceptions of beauty will be mistaken. A prominent

⁵⁴ This condition may seem to pose a challenge with natural kind categories that are highly sexually dimorphic. For example, the male mallard duck has a bright green head, whereas the female is a muted brown. In such cases we may meet the challenge by positing that we form two separate perceptual categories, one for the male and the other for the female. Or, we may posit that the green head feature is included as a relevant feature, but weighted less heavily in terms of diagnosticity.

example of this might be North American beauty pageants, where judges might be biased towards Caucasian beauty standards due to the predominance of Caucasians in the population. Insofar as the judges are attempting to judge beauty in the female face and form (as opposed to the female Caucasian face and form), then their perceptions and judgments of beauty will be in error.

All together, these are the basic elements of Waltonian perceptualism. However, there are many issues yet to be resolved. In what follows I try to resolve outstanding issues as well as to show how this account meets the desiderata laid out at the beginning of this chapter.

4.6 Meeting the desiderata

4.6.1 Cultural desideratum

Waltonian perceptualism has the virtue of being able to explain how beauty ideals change over time: they change as the relevant population statistics for a given category changes. It also explains how beauty ideals differ between cultures, as population statistics will differ in different cultures. Finally, it predicts that those factors that enhance processing fluency that do not require the ability to perceptually categorize an object will be common across all cultures, and so should lead to positive judgments regardless of culture. This final point can help explain why we are often intrigued by objects that belong to art practices we are not familiar with – such category-independent properties can serve as points of entry into learning about new categories of art.

4.6.2 Diversity desideratum

Waltonian perceptualism easily handles the diversity desideratum, which is the call to explain why it is that we find such a motley crew of objects beautiful, and what it is that unites all these things in absence of a general rule (and beyond the truism that they are all beautiful). Understanding the detection (or illusion) of beauty in terms of the fluency of perceptual processing provides the answer. We can perceive just about any kind of object to be beautiful, because the perceptual experience of beauty occurs via the relative fluency with which any kind of object is processed. The multiple factors that enhance or detract from processing fluency provide the answer to what unites this diverse group of objects.

This account also predicts that illusory perceptions of beauty will be commonplace because most of us are not perceptual experts with respect to many categories, and so this account also has the virtue of explaining the ubiquity of (erroneous) aesthetic judgments in those who are not perceptual experts with respect to a given category.

4.6.3 Abstracta desideratum

The abstracta desideratum calls for a theory of aesthetic perception to explain our practice of attributing aesthetic properties to non-perceptible objects such as theories, math proofs and conceptual art. Waltonian perceptualism can answer this call by pointing to non-perceptual processing fluency (Winkielman et al. 2003). The fluency with which cognitive processing occurs might give rise to similar aesthetic responses, though they would not be perceptions in this case. The fluency of cognitive processing is likely implicated in producing the

epistemic emotions, which are emotions that are produced in response to our cognitive predicaments (Morton 2010). They include emotions such as surprise, uncertainty, curiosity, and felt certainty. For example, we might reason to a conclusion that surprises us (perhaps fluent processing that abruptly becomes disfluent), or that we may feel uncertain about with respect to its truth (disfluent processing). In turn, these emotions are likely relevant in our aesthetic evaluations of non-perceptible objects. Perhaps the beauty of a math proof lies in the ease of understanding its progression.

Exactly how the details of this account should be expanded remain to be worked out – the proposal here is speculative. But the general form of the solution suggests that aesthetic properties can be part of perceptual experience, but also part of experience in a broader sense, where this is understood as part (or perhaps wholly constitutive) of cognitive phenomenology. If this is right then it may go a ways towards explaining the tendency of philosophers to retreat from aesthetic perception to a more broadly experiential view. There are genuine cases of non-perceptual experience of aesthetic properties. The existence of these cases, however, does not undermine aesthetic perception. There is no contradiction in holding a view whereby some aesthetic properties are perceived and others are not, or where aesthetic properties can be perceptible or part of cognitive phenomenology, depending on the case.

An alternate solution is to posit that attributions of aesthetic properties to these non-perceptible objects are figurative rather than literal. Perhaps aesthetic perception provides the basis for more metaphorical applications of those properties to non-perceptual works (Sibley 1959). When we speak of the plot of a novel as ‘unbalanced,’ ‘clunky,’ or ‘graceful,’ the paradigm instances we are drawing on that ground our understanding of this concept are

perceptual (Lakoff and Johnson 2008). Aesthetic perception is what allows us to make the metaphorical leap to aesthetically appreciate non-perceptual objects.

4.6.4 Affect desideratum

The affect desideratum is that a view of aesthetic perception should illuminate how it is that aesthetic properties can be both evaluative and perceived. On the fluency account that is part of Waltonian perceptualism, the fluency pertains to perceptual processing itself. The affect is produced as part of – or perhaps a byproduct of – producing perceptual experience. As noted previously, affect or the emotions have been proposed by several philosophers as a way of explaining the evaluative nature of aesthetic properties. Affect is positively or negatively valenced, and so does not just serve to describe, but also to evaluate.

Against this, one might worry that the affect, since it is a product of perceptual processing, is better said to be a post-perceptual experience after all, perhaps a component in one's non-perceptual phenomenology. I think this is the wrong view for the following reason. Processing fluency occurs as the perceptual system attempts to or actually makes a perceptual categorization. It is concurrent with, not just the result of, perceptual processing. So if we allow that the resulting non-aesthetic high-level property ('impressionist painting,' 'dog') is represented in perceptual experience, then we should allow that the aesthetic property is also represented. There are also additional reasons to think that some affect can be 'bound' into perceptual experience, but review of these would take me far afield of the main focus of this chapter (see for example Barrett and Bar 2009).

Another source of dissent might stem from the worry that this is too simple a view to capture the complex evaluations we make of artworks. We might, for example, evaluate a highly prototypical artwork that is thereby fluently perceptually processed to be derivative. This *prima facie* seems to be at odds with the predictions made by fluency theory, which predicts it will be judged as more beautiful. However, this worry can be addressed in two ways. First, we should distinguish between the perception of aesthetic properties and overall aesthetic evaluations (what Sibley called ‘verdicts’). Perceiving an object to be beautiful is compatible with an overall evaluation or verdict of it as derivative. Second, in a related point, we may posit that only some aesthetic properties are perceptible. It may be the case that we cannot perceive objects to be derivative, but only judge them to be so based on the aesthetic properties we do perceive, along with background knowledge of when the artwork was produced. Aesthetic perception need not entail that all aesthetic properties need be perceptible. The route from the perception of aesthetic properties to an overall evaluation may involve more than simply relying on our aesthetic perceptions – we may also draw on our background knowledge in order to arrive at our more considered evaluations.

4.6.5 Training desideratum

The training desideratum requires any account of aesthetic perception to provide a convincing story about how it is that training and the perception of aesthetic properties are compatible. Waltonian perceptualism resolves this by understanding training in terms of perceptual learning. On this account aesthetic training is largely perceptual training with the goal

of the construction of perceptual prototypes for new categories of objects, which can then allow for more or less fluent processing of individual category members.

Moreover, it explains why the sorts of categories that seem to require training are more likely to be categories of art. For natural objects such as dogs and horses, we have already undergone perceptual learning through the prevalence of these objects in our environments. Artworks, however, require more explicit training because of their relative absence from our developmental environments – children are typically exposed to relatively few exemplars of ‘impressionist paintings’ and ‘Dogon masks.’

4.6.6 Un-trained desideratum

The un-trained desideratum is to explain how some perceptions of beauty seemingly do not require training, while maintaining a unity between our trained and untrained aesthetic perceptions. Waltonian perceptualism meets both of these points. It provides three explanations for why some perceptions of beauty do not require training. First, the perceptual categories that we already possess – either through perceptual learning, or perhaps where these categories are innate – allow us to have perceptual experiences of beauty of many objects without formal training. For example, we likely acquire the perceptual category ‘dog’ through perceptual learning, and so can perceive dogs as beautiful or ugly. Second, given that some fluency-enhancing properties do not rely on experience with a category, then these properties will not require training. And third, some of these ‘perceptions’ of beauty will be illusory, so minimum proficiency with a category (where this is understood as the formation of a perceptual category) will result in such ‘perceptions.’

The account also provides a unified explanation for both trained and untrained aesthetic perception. Perceptual learning of categories can take place via passive exposure, but can also be accelerated via training. So it predicts that we should be able to learn to perceive objects as belonging to categories of art (or other non-art categories) through either method.

This unified explanation for our perception of aesthetic properties should be preferred to the explanation provided by the cognitive permeation interpretation, which must either (implausibly) hold that all aesthetic perception occurs via cognitive permeation, or grant that we can perceive some aesthetic properties without cognitive permeation, where this ability is left unexplained beyond positing a brute supervenience relation. So Waltonian perceptualism is more explanatory in this respect.

4.7 Outstanding issues

4.7.1 What is the correct category for viewing objects in?

Typically there will not be just one correct category – most objects will belong to several different categories. For example, a face might be beautiful as a face, a female face, as a face of a specific ethnicity. These categories are all correct. But of course, there are incorrect categories, such as male face, Caucasian face, and so on in which one can judge an object. So it is not a case that just anything goes. It does mean, however, that something might be beautiful with respect to one category and not with respect to another.

This is a constrained relativism – where the ascription of aesthetic properties is category-relative (‘beautiful for an F’) – and tracks common sense. We make judgments like these all the

time: ‘that’s the best of Picasso’s blue period but really not so remarkable in light of his total oeuvre.’

4.7.2 Are aesthetic judgments that do not rely on categories still possible?

There will be some aesthetic judgments that are not relative to a category because some fluency-enhancing features of objects do not depend on our perceptual expertise. However, despite the existence of such category-independent aesthetic properties, what we really care about in aesthetics are typically the category-dependent ones. When we strive to become aesthetic experts, it is always with respect to a given category, where this is usually rather narrow (Lopes 2018).

This is because there is no real expertise to be achieved with respect to category-independent aesthetic judgments. These (veridical) aesthetic perceptions are open to all, without training. Perhaps category-independent experiences of aesthetic properties serve as the hook by which we become interested in different categories of objects, and so are incentivized to specialize and develop expertise. I may not know, on my first encounter, whether Dogon masks are to be appreciated as a category of art or as another kind of artifact, but their category-independent beauty may intrigue me and cause me to investigate further.

4.7.3 Is there a limit to which categories we can become perceptual experts in?

Yes. Sometimes our categories will inevitably be biased, because of the limits of what we can be exposed to. Unlike machine learning, which can be trained on millions of exemplars,

we will likely form our prototypes based on a mere few thousand exemplars. So we are more likely to be biased in our prototypes because our sample is smaller. And of course we may be exposed to biased samples (though this is a problem in machine learning as well), with no practical way of remedying the issue. So these two issues can prevent us from becoming experts in a given category, because our perceptual categories will in turn be biased.

Also, there are likely some categories for which we cannot develop perceptual expertise because they are not even candidates for perceptibility. These categories will likely be highly abstract, without any or enough perceptual features to use to make reliable categorizations. It is an interesting question as to whether we can provide some principled guidance as to where the limits are, though one that I cannot answer here.

4.7.4 The naked mole rat problem

Naked mole rats are tiny hairless creatures that really appear to be quite ugly, regardless of how ‘prototypical’ they are qua naked mole rats. This objection assumes that those who have perceptual expertise with naked mole rats would nevertheless find all exemplars ugly, regardless of how prototypical. This assumption may itself be false. In Saudi Arabia there is an annual beauty contest for camels, where some of the criteria for camel beauty include having delicate ears and a fulsome snout. Competition is so fierce that some camels undergo plastic surgery (Weaver 2018). While camels are no naked mole rats, outsiders without expertise might look upon the category of camel as a whole and find it aesthetically unimpressive – they are rather strange and gangly creatures with their spindly legs and bulky humps. Nevertheless, the presence of camel beauty contests suggests that once some expertise is developed, then people do really

come to perceive some camels as beautiful. Perhaps it is the same with naked mole rats – the answer to this question awaits empirical scrutiny (there are no naked mole rat beauty contests that I am aware of).

Suppose it is not the same however, and it turns out that naked mole rats are just ugly. In this case we might point out that just as there are fluency-enhancing properties that are independent of category, there are also such fluency-impeding properties. If all naked mole rats possess some of these properties, then this may explain our general perceptions of them as ugly without expertise. If this is the case, then even if someone does develop expertise with naked mole rats, they may not come to perceive certain instances as beautiful.

4.7.5 The Cindy Crawford mole problem

Cindy Crawford was a supermodel from the mid-1990s. She had a large, dark mole on the upper right side of her lip. This mole was asymmetric, and yet she was considered one of the most beautiful women in the world at the time. How to explain this?

Crawford possessed many of the other category-independent features that enhance fluency processing. These alone may explain judging her as beautiful, despite the mole. However, we might instead posit that the mole allows us to identify Crawford more rapidly, so enhances fluent perceptual processing. Overall, this points to the need for further work to understand how aesthetic ‘flaws’ can sometimes enhance our aesthetic judgments. It thus remains to be seen whether aesthetic flaws can be given an explanation in terms of processing fluency. The Japanese concept of ‘wabi-sabi’ – which celebrates imperfections as adding to beauty – is perhaps a good place to begin investigation.

4.7.6 How can this account be extended to other aesthetic properties?

Here one might worry that fluency is too blunt an instrument to explain the wide range of aesthetic properties. The account seems to work most straightforwardly with beauty and ugliness, and for all I have said here, it is unclear that it applies to all or even some other aesthetic properties. There is not yet any empirical work we can point to that provides guidance in this area.

If the account were only able to establish that beauty and ugliness are represented in perceptual experience, then this already would be a considerable achievement, given the central position in our lives of these aesthetic properties. Nevertheless, it would be premature to rule out the possibility that fluency has a role to play in explaining how some other aesthetic properties are perceived.

I do not have a fully worked out answer, but only some rather speculative suggestions for how this might go. How we represent further aesthetic properties in perceptual experience might be understood by making more fine-grained distinctions with respect to the different components of perceptual categories that lead to more or less fluent processing. For example, perhaps some aesthetic properties can be understood as more specific ways in which something can be beautiful or ugly. Aesthetic properties such as graceful/clumsy, harmony/disharmony, and tasteful/garish might all be cases that are explicable in terms not just of relative (dis) fluency but also in terms of which properties are standard or contrastandard. That is, perhaps our understanding of fluency needs to be broken down into the ways in which fluency can be enhanced or hindered, or by the reference to the specific properties of the object that cause the

enhancement or hindering. If we understand gracefulness as movement or form that follows a predictable pattern without deviations, then this might explain not only how fluency is enhanced but also characterize gracefulness (though particularist considerations might count against providing any such unified characterization).

We might also draw on Walton's suggestion that contrastandard properties also influence what aesthetic properties are perceived. A twitching kinetic painting, he suggests, may be perceived as shocking, because it is contrastrandard for paintings to possess this property. Again, we might gain more understanding not just by pointing out that this contrastandard property impedes fluent categorization, but in analyzing the way that it does so, or by referring to the specific properties in virtue of which it does so. A twitching painting might be shocking, a bleeding painting might be gruesome.

Understanding perceptual categories to have a prototypical structure also gives us a wide range of variables to draw upon. Typical features diagnostic of category membership can be understood as varying along a range, for example (this is called a 'dimension'). So the feature of typical size may in fact be represented as a range of sizes, with a cut-off for smallest and largest. We might then begin to understand some aesthetic properties as emerging out of where exemplars fall on this spectrum. For example, cute things might all be at the extreme limit for possession of standard properties (they might be the smallest objects in the permitted size range). In the case of cats, for example, kittens may be perceived as cute in virtue of their possessing a very small body and large head and very wide-set eyes. Why this should be a (generally) positively valenced property while nevertheless existing at the category margins is a question for further investigation.

4.8 Conclusion

In conclusion, in this chapter I have proposed a novel view of how we come to perceive beauty that explains how (at least some) aesthetic properties can be both learned and perceptual. On Waltonian Perceptualism, perceptual learning explains how we come to represent new categories of objects in perception, the fluency hypothesis explains how these categories lead to perceptual experiences of beauty and ugliness (and perhaps aesthetic properties more generally), and an understanding of what it is to be a perceptual expert allows us to say when those experiences are veridical. This account has the virtue of being able to satisfy several desiderata for a theory of aesthetic perception. While there are several outstanding issues that the account must resolve, it is nevertheless a more complete account of aesthetic perception than that of cognitive permeation.⁵⁵

⁵⁵ My gratitude to Dominic McIver Lopes, Christopher Mole and Elisabeth Schellekens for comments on an earlier version of this chapter, as well as comments from Kendall Walton, Stacie Friend and David Davies on a short article forthcoming in *JAAC* (2020) that contains some of the themes I elaborate on here. Thanks also to the audiences of 2018 Annual ASA meeting in Toronto, attendees of the 2019 Camp Aesthetics in Capitol Reef National Park in Utah, and of a 2018 work in progress seminar at the University of British Columbia.

Chapter 5: The Epistemology of Perceptual Learning

Summary: Perceptual learning involves long-term changes in the contents of perception - experts and novices have different perceptual experiences, and the perceptual experiences of experts are richer in some respects. This suggests that experts may form a greater variety of immediately justified perceptual beliefs as compared to novices. However, this view faces a challenge: if perceptual learning is a form of cognitive permeation, then it is not a source of immediate justification. Either the justification our learned perceptual experience provides is mediate, or it fails to provide any justification at all to the resulting beliefs. Against this challenge I argue that perceptual learning does not require cognitive guidance. In cases where cognitive guidance is involved, it does not suffice for cognitive permeation.

5.1 Introduction

A novice and a wine connoisseur are drinking some wine together at a party. The connoisseur recognizes the wine to be a zinfandel. The novice does not.

An ornithologist and a novice are out in the forest. They both recognize a bird in a tree. The ornithologist recognizes it as a barn owl. The novice does not.

A dog fancier and her novice friend are walking in the park. They both recognize a dog as it runs by. The dog fancier recognizes it to be a Boston terrier. The novice does not.

In all of these examples, the expert is able to form beliefs that the novice is not able to. What is the nature of the expert's ability? Some might claim that the novice and expert both have identical perceptual experiences, but that the expert is able to do more with her perceptual experience. She is able to make more inferences on the basis of her perceptual experience, or attend to key components of her experience over others, thanks to the conceptual tools at her disposal due to training (for example Pylyshyn 1999).

I have argued in chapter three against this view of expertise; the learned attentional patterns of perceptual experts are themselves partially constitutive of a change in perceptual experience. Henceforth I will assume that the expert's perceptual experience is different from the novice, and that this perceptual difference is what explains their differing abilities. We can further specify the nature of this perceptual difference by appealing to a debate in the philosophy of perception over what sorts of contents perceptual experiences can have – the content debate.

5.1.1 The content debate

To say that perceptual experiences have contents is to say that such experiences have accuracy conditions – they convey to the perceiver that the world is a certain way, where what is conveyed can either be accurate or inaccurate. These contentful experiences are sometimes understood in representational terms, where representations are mental objects with semantic properties such as contents, reference and truth conditions. The issue of interest is then what sorts of properties we can come to represent in perceptual experience.⁵⁶

Sparse content theorists hold that we can only represent low-level properties in our perceptual experience, where low-level properties are confined to the most basic building blocks needed to construct perceptual experience, such as colours, illumination, shapes, motion and spatial properties (Tye 1995; 2000; Dretske 1981; Clark 2000; Raftopoulos 2009; Connolly 2014; 2019). *Rich content* theorists endorse the rich content thesis: we represent some high-level properties in some of our perceptual experiences, where high-level properties are simply all those properties that fall outside the scope of low-level properties (Siegel 2010; 2006b; Fish 2013; Werner 2016; Siewert 1998; Brogaard 2018; Stokes 2018a). Such theorists vary in terms of which high-level properties they take to be perceptually represented, but these may include natural kind properties, artifactual kind properties, and aesthetic and moral properties.

While others have argued that the perceptual experiences of experts differ from those of novices only with respect to low-level properties (Connolly 2019; 2014), here I adopt the view

⁵⁶ See (Siegel 2010 ch. 1&2) for an extensive defence of the thesis that experiences have representational contents.

that they also differ with respect to high-level properties. Experts can come to represent high-level properties such as ‘barn owl,’ ‘zinfandel,’ and ‘Boston terrier’ in their perceptual experience, whereas these high-level properties are not perceptually represented in the experience of the novice. For arguments in favour of this view, see chapter three. I understand his ability of experts as one that involves perceptual categorizations – categorizations that are performed by the perceptual system itself, and that allow for the attribution and perceptual representation of high-level properties. Perceptually categorizing something as a dog allows one to represent the property ‘dog’ in perceptual experience.

Acquiring the ability to make reliable perceptual categorizations is undoubtedly not the only requirement for becoming an expert. But it is sufficient for one component – perceptual expertise. Perceptual expertise as I will use the term here refers to the ability to come to represent high-level properties in perceptual experience, where this is accomplished through developing a perceptual category for the property. It is thus broader than the way it is used in the psychological literature, insofar as psychologists that study perceptual expertise are usually interested in a small subset of high-level properties – those termed ‘subordinate categories’ by Eleanor Rosch and colleagues (Rosch et al. 1976).

Subordinate categories are the least inclusive category of those in a three -level hierarchy composed of subordinate, basic-level and superordinate categories.⁵⁷ Subordinate categories include zinfandel, barn owl, and Boston terrier. Basic-level categories include wine, bird, and dog. Superordinate categories include animal and liquid. Studying acquired subordinate categories is relatively easy because few of us are perceptual experts in these domains. This is in

⁵⁷ More levels can be added to this hierarchy, but for simplicity’s sake I discuss only the original three proposed by (Rosch et al. 1976).

contrast to basic-level categories, where most of us are experts at perceptually categorizing things as dogs and birds and the like. While the empirical study of perceptual expertise with subordinate categories is an important source of evidence for the view I have adopted here, perceptual categorization should be understood as potentially involving superordinate and basic-level categories as well.

In summary, the difference between novice and expert is one in which experts can come to represent high-level properties in their perceptual experience that novices cannot, where these high-level properties are understood along the lines of subordinate, basic-level or superordinate categories. In the examples provided above, the ‘novices’ are nevertheless perceptual experts with respect to the high-level properties ‘dog,’ ‘wine,’ and ‘bird.’ The question that will concern me in the rest of this chapter is the epistemic status of the expert’s perceptual beliefs.

5.1.2 Immediate justification

Perception is often thought to provide a distinctive source of justification for our beliefs, with some maintaining that the beliefs we form on the basis of our perceptual experience are basic, or immediately justified (Huemer 2001; Brewer 1999; Pryor 2000; Goldman 2008; Kornblith 2002; Markie 2006). A belief is immediately justified when it does not depend on another belief for justification, but is nevertheless justified (Pryor 2005; Alston 1983; McGrath 2017). Mediate justified beliefs are those that do depend on one or more other beliefs for justification. In the case of beliefs formed on the basis of perception (henceforth post-perceptual beliefs), the source of justification is not a further belief, but rather our perceptual experience.

Post-perceptual beliefs are therefore plausibly immediately justified (Wright 2002; 2004; though see Kvanvig and Riggs 1992).⁵⁸

If the number and kinds of properties we represent in perceptual experience are fixed, then so too is the scope of immediate perceptual justification. However, if our perceptual experience can change and grow with learning, then it seems that the scope of immediate perceptual justification can also change. Call this the *wider scope* thesis: the scope of immediately justified perceptual beliefs can be enlarged or broadened through learning. This change in scope would have particular significance for foundationalism, the view that mediately justified beliefs are ultimately ‘anchored’ in immediately justified beliefs (for example Bonjour 1978; 1985). Beyond foundationalism, any view that allows that perception provides us with immediately justified beliefs should care about whether the wider scope thesis is true.

At first glance, the view of perceptual expertise that I have endorsed above would seem to unambiguously support the wider scope thesis: perceptual experts come to have richer perceptual experiences, and since perceptual experience is a source of immediate justification, then their resulting beliefs are also immediately justified. In fact, some foundationalists do appeal to learning as a way of increasing the number of basic beliefs one has (Goldman 2008; Brewer 1999). Likewise, rich content theorists tend to hold that many high-level properties come to be represented in perception via some form of learning (Macpherson 2012; Stokes 2018b; 2020; Siegel 2010; 2006b). While our perceptual systems may allow us to represent certain low-level properties like colours and shapes or limited high-level properties such as causation without any

⁵⁸ Those who endorse the claim that perception can provide us with immediately justified beliefs tend to endorse *dogmatism or phenomenal conservatism*. As (Huemer 2001, 99) characterizes it: “If it seems to S as if P, then S thereby has at least prima facie justification for believing that P.” See also (Pryor 2000; Tucker 2010). However, one can endorse immediate justification without endorsing dogmatism, and so I focus here on the former.

learning (Siegel 2009), the idea that the visual system represents high-level properties like Boston terriers and barn owls from the get go is less plausible. Some of us recognize these objects, and some of us do not – learning plausibly explains the difference.

In both debates, however, the nature of the learning process must be further specified so that the epistemic consequences can be better understood. Understanding what sort of learning is at issue is particularly important because the wider scope thesis and the rich content thesis can come apart – we might represent high-level properties in perceptual experience without its being the case that this provides a new source of immediate justification.

In what follows I explore a way in which the epistemic status of expert perceptual beliefs may be undermined: they may be the result of cognitive permeation. In such a case, the resulting perceptual beliefs may be mediately justified by the perceptual experience or receive no further justification at all from the resulting experience. In section 5.2 I provide an overview of cognitive permeation, how it might be used to explain perceptual expertise, and what the epistemic implications of this view are. In section 5.3 I argue that perceptual expertise is not the result of cognitive permeation: cognition plays no necessary role in perceptual learning for expertise, and the sorts of concepts involved are proprietary to the perceptual system.

5.2 Cognitive permeation

Cognitive permeation, also commonly referred to as cognitive penetration, is hypothesized to occur when the contents of perceptual experience are altered in some way by one's cognitive states, such as beliefs and desires (Pylyshyn 1999; Macpherson 2012; Stokes 2013; 2018b;

Briscoe 2015; Raftopoulos 2009; Newen and Vetter 2017; Gross 2017; Siegel 2012; 2013).⁵⁹ In the discussion that follows I will focus on beliefs because if perceptual expertise is the result of cognitive permeation, then the most plausible cognitive states involved will be beliefs (I will also discuss concepts in section three). These permeating beliefs are sometimes called ‘background beliefs.’ Siegel (2012) offers the example of Jill, who has a background belief that Jack is angry with her, and on this basis has the perceptual experience of Jack’s face as angry. This experience in turn causes her to form the post-perceptual belief that Jack is angry with her, or reinforces her existing belief.

The basic idea of cognitive permeation is easy enough to grasp. Providing an adequate characterization of the phenomenon, however, has proven more challenging. There is no wholly uncontroversial definition (Stokes 2013; Zeimbekis and Raftopoulos 2015b, 27–32).

Nevertheless there are common strands amongst several definitions.

5.2.1 Defining cognitive permeation

First, there is a *causal* condition: the background belief must be the cause of the resulting perceptual experience. Second, there is an *internal* condition: in order for cognitive permeation to occur, the background belief must provide direct input to perceptual processing, which then modulates perceptual experience. Third, there is a *semantic coherence* condition: the content of the input provided by the background belief must bear a semantically non-arbitrary relationship to the resulting content of the perceptual experience. That is, the influence of the cognitive input

⁵⁹ Though traditionally the term is ‘cognitive penetration,’ I follow Becco Copenhaver (personal conversation) in using the term ‘cognitive permeation’ as it paints a less gendered picture of the relationship between perception and cognition.

on the contents of perception must be “coherent or quasi-rational when the meaning of the representation is taken into account” (Pylyshyn 1999, 365 n. 3).⁶⁰

Together, these conditions suffice for cognitive permeation, and rule out common cases of benign cognitive influence on perception. For example, suppose that I believe my friend is walking through the door and so turn to look at the door, but I find the doorway to be empty. This does not count as cognitive permeation, even though my belief has altered the contents of my perceptual experience. While my belief (along with a concomitant desire) has caused me to change what I am looking at, this effect is trivial insofar as there is no semantic relation between the contents of my belief and the contents of my perceptual experience. Our beliefs and desires exert a near-constant effect on how we orient ourselves with respect to the world, but this on its own does not count as cognitive permeation.

However, now suppose that it turns out that luckily my friend happens to be walking through the doorway as my belief causes me to turn and look. Here, there is a semantic relation between my belief (‘my friend is in the doorway’) and the contents of my perceptual experience (something along the lines of ‘my friend is in the doorway’). Nevertheless, there is no internal link between the belief and the contents of my perception. Instead, the contents of perceptual experience were altered via an indirect causal link: my belief caused a reorientation of my attention to a particular region of space, which then altered the input to my perceptual processing.

Given that an internal link is difficult to ascertain, a counterfactual causal analysis of cognitive permeation is helpful to ruling out such cases of ‘lucky’ semantic coherence from

⁶⁰ While some theorists have rejected the semantic coherence condition in order to make room for the permeation of experience by desires and moods, the sort of cognitive permeation hypothesized to occur with expertise must be explicable in terms of beliefs or concepts closely tied to the resulting perceptual content, and so would meet the semantic coherence condition.

counting as cognitive permeation. Siegel suggests that the following counterfactual holds in most cases of cognitive permeation: “If the subject were not in background state B but was seeing and attending to the same distal stimuli, she would not have an experience with content p” (Siegel 2012, 6). Here ‘background state’ is meant to refer to whatever background belief, desire or other cognitive state that is permeating perceptual experience. The counterfactual will not hold for all cases because there may be instances where an alternate background state B* also causes an experience with content p, say. (Or, in a more far-fetched counterexample, an experience with content p might be produced by a whack to the head or divine intervention.) But this analysis is nevertheless useful. For example, when I believe my friend is walking through the doorway and she happens to be there, this fails to satisfy the counterfactual because if I failed to possess the background belief B but had turned to look at the doorway for some other reason, I would still have an experience of her in the doorway.

5.2.2 Diachronic cognitive permeation

Perceptual expertise as presented above seems quite different from cognitive permeation – psychological studies have found it requires long periods of training involving perceptual exemplars for a change in perceptual abilities to occur, not the simple acquisition of beliefs (Gauthier et al. 1998). No matter how firmly a novice might believe based on the testimony of an expert that the bird pictured is a barn owl, or that the wine she is tasting is a zinfandel, this does not alter her ability to recognize these things – acquiring such knowledge is not sufficient for perceptual expertise.

Nevertheless, some have argued that perceptual expertise should count as a form of

cognitive permeation: *diachronic* cognitive permeation, where this signifies that the permeation of perception by cognition happens over time, with cognition gradually reshaping the structure of the perceptual system and only eventually leading to a change in perceptual experience (though see Arstila 2016; Raftopoulos 2001a; Churchland 1988; Pylyshyn 1999). This is unlike the case of *synchronic* cognitive permeation described above, where beliefs are immediately able to permeate experience.

Indeed, as Paul Churchland (1988, 175) points out, the original proponents of cognitive permeation take pains to emphasize that it occurs over extended periods of time as the result of training and practice (Churchland 1979 ch. 2; Kuhn 1962 ch. V, X). This provides a greater fit with perceptual expertise as outlined above. Churchland himself offers the example of musical expertise, where the expert “perceives, in any composition whether great or mundane, a structure, development, and rationale that is lost on the untrained ear” (1988, 179). Other more recent defenders of cognitive permeation have put forward similar cases of perceptual expertise such as the radiologist’s expertise at detecting tumours (Siegel 2012; Stokes 2020), the art connoisseur’s ability to recognize categories of artworks (Stokes 2014) and the herpetologist’s ability to recognize copperhead snakes (Lyons 2011)⁶¹.

If perceptual expertise is a form of diachronic cognitive permeation, then it must be the case that it is the learning of the relevant theory, or some subset of these beliefs, that is responsible for altering one’s perception (Fodor 1988, 195). Given that these beliefs need not be responsible for one’s perceptual experience in the moment, but instead exert their effect over

⁶¹ See (Cecchi 2014) for an argument that perceptual learning of low-level properties is the result of diachronic cognitive permeation via top-down attention. I will not address this argument further here, but it does not seem tenable given the abundance of studies on low-level perceptual learning that demonstrate that it can occur in the absence of top-down attention, or even absent conscious perception (for example Seitz and Watanabe 2009; Watanabe, Nanez, and Sasaki 2010).

time, they are subtly different from the background beliefs involved in synchronic cognitive permeation. I will nevertheless continue to call these beliefs responsible for the restructuring of the perceptual system during the training period ‘background beliefs’ in the interest of avoiding an over-proliferation of terminology.

Exactly how to fill in the details of the relationship between background beliefs and the subsequent perceptual experience requires some work. We can again appeal to the conditions for cognitive permeation outlined above. First, the background beliefs must cause the change in the perceptual experience. Second, the content of the background beliefs must bear a semantically non-arbitrary relationship to the resulting content of the perceptual experience. Given that we are interested here in high-level properties that correspond to categories, the background beliefs will plausibly be beliefs about the nature of these categories. Novices may begin by acquiring background beliefs roughly of the form ‘objects of category C possess properties x, y, z’ where x, y, z are low-level perceptual properties such as size, colour, and shape, and C corresponds to a high-level property such as being a barn owl. The end result is then one of the novice becoming a perceptual expert, and thereby representing the corresponding high-level property in her perceptual experience.

The third condition, that there must be an internal connection between the belief and the perceptual experience, is more difficult to adopt here. It is hard to see how the causal link between perception and cognition could be direct and internal and yet cognition only exert its influence on perceptual processing, and perceptual experience, over time. That is, if perceptual experience is not synchronically cognitively permeable by the learned background beliefs, then how could it ever come to be diachronically cognitively penetrable? A gradual effect, if it is to be internal, still requires a direct pathway for providing input.

5.2.2.1 Diachronic cognitive permeation via attention

One plausible story for how diachronic cognitive permeation occurs involves appealing to attention. This may be confusing at first. The discussion above seems to rule out attention. One reason for ruling out attention is on the basis of its violating the internal condition because attention ostensibly serves only as an indirect modulator of perceptual experience. Another reason for ruling it out is on the basis of the semantic coherence condition because attention ostensibly determines only the input to perceptual experience and not the resulting content. However, several theorists have recently argued that some varieties of attention should be understood as a source of cognitive permeation (Mole 2015; W. Wu 2017; Stokes 2018b).

Explaining how attention might allow for diachronic cognitive permeation requires making some more fine-grained distinctions between different varieties of attentional phenomena (Carrasco 2011; K. M. Armstrong 2011; Posner 1980; Findlay and Gilchrist 2003; Ransom, Fazelpour, and Mole 2017). Most relevant to perceptual learning are the distinctions between spatial and object or feature-based attention, exogenous and endogenous attention, and overt and covert attention. Our attention can be oriented on the basis of spatial location, features or objects. Spatial attention is attention to a location, irrespective of what is there. Object and feature-based attention are attention to objects or features, irrespective of their location. Exogenous attention is a ‘capturing’ of our attention by certain features, irrespective of our aims: we involuntarily turn to look at sudden loud bangs and bright flashes. Endogenous attention is driven by cognition, where this is usually understood in terms of our goals and tasks. We attend only to people wearing red shirts in a crowd because we are looking for a friend wearing a red shirt; we attend

to a given spatial location on the basis of a friend's pointing, though we do not know what we will see. Overt attention occurs when there is an attentional shift that involves a reorientation of one's body, where this includes eye movements. Covert attention occurs when there is a shift in attention without any such bodily change, such as when we tune out the person we are conversing with in order to eavesdrop on the conversation at the table next to us. All of these distinctions are orthogonal to one another. For example, covert attention can be endogenous or exogenous, and can involve features, objects or spatial locations.

Most relevant for our purposes here, Chris Mole (2015) argues by appealing to a body of psychological research that perceptual experience is altered by our learned concepts through covert attention (though see Gatzia and Brogaard 2017). He appeals to evidence that our acquired number and letter concepts influence the way attention is allocated to a visual array. On an influential theory of attention, it is understood in terms of biased competition for the brain's representational resources, where processing at multiple levels (including perceptual processing) can be biased by several factors in favour of certain features or spatial regions over others, and the winners of this competition become those items or regions that are attended (Desimone 1998; Duncan 1998). If we take on board this theory, then the evidence Mole appeals to suggests that such concepts are able to bias the competition during perceptual processing itself: concept-driven attention affects the way perceptual processing occurs, and does not play a mere orienting role as with cases of endogenous spatial attention.

While Mole focuses on cases where expertise with the relevant category has already been established, we may elaborate the picture to include an account of how diachronic cognitive permeation may occur – call this the *belief-driven hypothesis*. First, learning of the background beliefs occurs: a person learns via testimony that a given object possesses characteristic features

x, y, z (‘robins have red breasts, brown feathers and a yellow ring on their beaks’). This step might be called concept acquisition or concept amplification.⁶² Second, on the basis of this learned concept the person overtly and endogenously directs her attention to these characteristic features one by one to make an identification. This is in part a post-perceptual process, likely effortful at first, whereby the person accumulates perceptual evidence for concept application. It is a process somewhat akin to checking boxes on a list (‘Does the bird have a yellow ring on its beak? Check.’) Third, over time repeated attention to these features causes the restructuring of the perceptual system: the pattern of attentional tuning becomes automated and covert, no longer requiring the deployment of endogenous or overt attention – it directly privileges the grouping of these features and renders them simultaneously salient, perhaps by biasing the competition for perceptual processing resources (Desimone 1998; Duncan 1998; Mole 2015). This automatic perceptual grouping may even be said to correspond to the representation of high-level properties in perceptual experience – the relevant low-level properties have been grouped in a way that reliably tracks the relevant natural or artifactual kinds (see also McGrath 2017 for discussion of how recognitional abilities may be developed over time).

Much of this account is right. Perceptual expertise does indeed involve learning to attend to the features relevant for object categorization, and the account I develop in chapter three appeals to attentional tuning over time – known in perceptual learning as ‘attentional weighting.’ However, the crucial point of disagreement here is whether background beliefs are necessary for perceptual expertise. In what follows I will argue that background beliefs are not necessary, and so the process by which we acquire perceptual expertise is not accurately considered a case of

⁶² Some might object here that concepts should not be understood in terms of beliefs. The challenge for those who nevertheless endorse cognitive permeation is then to articulate how theory or belief can lead to concept acquisition in a way that does not appeal to the pattern-driven hypothesis that I defend below.

cognitive permeation. First however, it is important to understand what is at stake. If perceptual expertise is the result of diachronic cognitive permeation – as many have assumed – then what is the epistemic status of the expert’s perceptual beliefs?

5.2.3 Epistemic implications of cognitive permeation

In some cases of cognitive permeation, the resulting perceptual experience seems to be unable to provide any justification at all to the resulting perceptual belief. Siegel’s case of Jill and Jack is arguably of this variety: given that Jill’s perceptual experience of Jack’s anger is caused entirely by her antecedent belief that Jack is angry, it seems illicit for the perceptual experience to provide any additional justification for the belief that Jack is angry. Siegel’s analysis of why this is the case rests on the notion of perceptual sensitivity to the incoming stimulus (2012, 4,12).⁶³ In such cases, the relevant content of the incoming stimulus (Jack’s actual facial expression) makes no difference at all to Jill’s perceptual experience. Instead, her experience is entirely dependent on her belief, in the sense that it – rather than the incoming perceptual information – determines her perceptual experience. Her perceptual experience would be of Jack as angry, regardless of the actual perceptual input (see also Lyons 2011, 301). Siegel holds that due to this insensitivity the resulting perceptual experience fails to provide any independent justification at all for the ensuing belief. With respect to this kind of cognitive permeation, while perceptual experience can come to have rich contents as a result of learning, the scope of immediately justified perceptual beliefs is not wider as a result.

⁶³ Though Siegel (2012) also appeals to epistemic circularity to explain the issue, I focus here on sensitivity. She has since provided an alternate analysis of hijacked experience in terms of poor inference (Siegel 2017).

However, this analysis is compatible with the possibility that there are cases of cognitive permeation that do allow for some sensitivity to the incoming stimulus – it is meant to explain why certain instances of cognitive permeation are epistemically problematic, not a comprehensive diagnosis of all cases of cognitive permeation. Indeed, Siegel (2012) and others (Lyons 2011; Stokes 2020) assume that in some instances cognitive permeation will be virtuous – such as the radiologist who can perceptually recognize tumours and the herpetologist who can recognize copperheads.

This assumption has received a limited defence. Lyons (2011) proposes that cognitively permeated perceptual experience is virtuous when it increases reliability. Stokes emphasizes the superior performance of perceptual experts in their domain of expertise – where this always involves cognition though in some cases perhaps not cognitive permeation – arguing that it is an epistemic good because “the expert is moving closer to an optimal cognitive stance on the world (or a part of it), where she can better acquire behaviourally relevant category and diagnostic information” (Stokes 2020 IV.3).

We can recast a defence in terms of sensitivity. It is plausible that the background beliefs of experts (of the variety we are interested in here) do help increase the sensitivity of the expert to the incoming stimulus rather than decrease it. For example, suppose my belief that dogs have tails, fur, and so on diachronically permeates my perceptual experience in order to allow me to identify something as a dog. This belief may allow me to become more sensitive to certain aspects of my experience, as I may preferentially attend to the animal’s tail and fur, say. And of course it allows me to become sensitive to the property of something’s being a dog. So long as the background belief does not cause me to perceptually categorize things as dogs irrespective of whether they are dogs, then it would seem that the belief is working in concert with perceptual

experience rather than ‘hijacking’ the experience or reducing sensitivity to perceptual inputs. If perceptual expertise results from diachronic cognitive permeation, it is plausible that perceptual experience nevertheless provides justification for the resulting post-perceptual beliefs.

Nevertheless there remains a challenge to the wider scope thesis on this view: perceptual experiences formed as the result of cognitive permeation may provide only mediate justification to the ensuing beliefs.⁶⁴ One reason for thinking this is because it may be reasonable to endorse what I will call the *dependence* principle, as articulated by (Silins 2013, 25): “When an experience has the content that p in part in virtue of some belief one has, it gives one justification to believe that p only in virtue of one’s having justification for the belief.” (McDowell 1982, 478; Plantinga 1993, 99–101; Markie 2005, 350; see also Cummins 1998).

One might resist the dependence claim on the grounds that justification and etiology are separate - how our perceptual experiences come to be is one thing, and whether they provide justification is another (Pryor 2000; 2005).⁶⁵ We may grant that background beliefs causally affect our perceptual experiences, but maintain this is a separate issue from whether the resulting perceptual experiences can grant immediate justification. In such cases, Pryor (2000, 540) claims that background beliefs are like putting on a pair of sunglasses – while they both causally affect our perceptual experience, it is not the case that they compromise immediate justification.

However, this response is not tenable in the face of the more extreme cases of cognitive permeation discussed earlier. Such cases arguably demonstrate that the etiology of perceptual experience does matter to whether or not perceptual experience can allow post-perceptual beliefs

⁶⁴ For alternate ways of motivating the claim that learned perception provides only mediate justification see (Vaassen 2016; McGrath 2017, 30–34).

⁶⁵ A different way of resisting the principle might draw on (Lyons 2011, 296), though he is not concerned with immediate justification. See also (Siegel 2017, 47–50), who considers various ways in which perceptual experiences may provide immediate justification even if they are cognitively permeated, and (Siegel 2017, 144–45) for discussion of expert judgments that result from cognitive permeation.

to be justified at all (Siegel 2012; 2017; Silins 2013, 25; Vance 2014; McGrath 2013). If this is right, then it seems reasonable to suppose that there may be cases where post-perceptual beliefs are only mediately justified by perceptual experience.

Sillins (2013) motivates the dependence principle by providing the example of Alexa, who forms an unjustified background belief that ‘things with such and such a look are sheep dogs,’ which then allows her to represent the property ‘sheep dog’ in perceptual experience. He appeals to the intuition that it would be illicit for Alexa to gain justification through her perceptual experience for the belief ‘that is a sheep dog’ given that she was not justified in the background belief that led to her experience in the first place. Those who share this intuition should endorse the dependence principle to explain why perceptual experience fails to result in an immediately justified belief. If we allow that unjustified background beliefs prevent the subsequent post-perceptual beliefs from being justified at all, then this would seem to indicate that the background beliefs are playing a mediating role.

Another consideration in favour of the dependence principle rests on my adoption of the semantic coherence condition. This condition ensures that there is a non-arbitrary semantic relation between the belief and the contents of perceptual experience, where this outstrips the sort of brute causal (non-semantic) relation that a pair of sunglasses has to perceptual experience. What is more, because the content of the background belief is connected in a quasi-logical manner to the content of the perceptual experience (even if this falls short of inference), this renders it plausible to also attribute a relationship of epistemic dependence. At minimum, those who want to resist the dependence claim must provide some positive reason to believe that etiology does not matter in the case of perceptual expertise, or does not impact immediate justification.

In summary, if it is true that expert perceptual experiences are the result of cognitive permeation, then while this may not render expert perceptual beliefs unjustified, it plausibly renders them only mediately justified. Therefore, perceptual expertise does not vindicate the wider scope thesis: experts do not broaden the scope of their basic perceptual beliefs through learning. This conclusion underscores the importance of gaining a proper understanding of the etiology of perceptual expertise. In the following section I argue that perceptual expertise is not the result of diachronic cognitive permeation. It is rather the result of perceptual learning, which is not necessarily cognition-involving.

5.3 Perceptual expertise is not the result of cognitive permeation

In the content debate, people have often assumed cognitive penetration as the default or most plausible option for explaining how learning can alter the contents of perceptual experience (though see Brogaard and Chomanski 2015). It is a tempting option because declarative learning – the learning of verbally conveyed facts – and belief formation so often occur before or coincide with the change in perceptual experience, so it seems natural to attribute the change to these newly acquired beliefs. For example, Paul Churchland, considering the hypothesis that the Müller-Lyer illusion is the result of learning, writes that if the hypothesis were true, then this would mean that “The illusion exists in the first place only because the relevant processing module is the well-trained victim of some substantial prior education – that is, of some penetration by cognitive activity” (1988, 174).⁶⁶

⁶⁶ For a more recent argument that the illusion is the result of learning see (McCauley and Henrich 2006).

Another example of this line of thinking can be found in Fiona Macpherson's writing: "People who subscribe to the existence of high-level content in visual experience are likely to reject cognitive impenetrability. This is because many of the arguments for high-level content proceed by arguing that *learning* can affect which visual experience one has." (2012, 32 emphasis original).

Jack Lyons writes: "Perceptual identification involves top-down penetration in that, among other things, one has to have learned what dogs and bananas look like" (2011, 291).

Finally, Robert Briscoe: "Hence, if psychological findings suggest that high-level properties *are* represented in the phenomenal contents of visual experience [...] then this would be seemingly good reason to suppose that visual phenomenal content has been penetrated by information originating outside of the visual system" (2015, 182).

However, this inferential leap is unwarranted. While appeals to cognitive permeation in cases of learning locate the source of the change in our acquired background beliefs, the direction of dependence may run the other way round: the changes to perceptual experience through repeated exposure to and practice with the perceptual stimuli are what themselves anchor and solidify the learning of new beliefs. That is, what is known as perceptual learning may instead be driving the changes in perceptual experience.

5.3.1 Perceptual learning

Perceptual learning has emerged more recently as an alternate way of explaining the sort of learning that takes place between novice and expert (Connolly 2014; 2019; Brogaard and Gatzia 2018; Brogaard 2018; Chudnoff 2018; Vaassen 2016; Arstila 2016). Perceptual learning may be

characterized as an enduring change in the perceptual system due to practice or repeated exposure to a perceptual stimulus. The perceptual system can be understood as signifying whatever cognitive resources are causally responsible for producing perceptual experience. As a result of perceptual learning, there is a long-term change in the contents of our perceptual experience (Goldstone 1998; Connolly 2019; E. J. Gibson and Pick 2000; J. Gibson and Gibson 1955a).

There is considerable evidence that with practice our perceptual experience can become more fine-grained and accurate in terms of low-level properties. For example, we can come to identify the direction of motion of small dots with better precision (Ball and Sekuler 1987; 1982) or the position of tiny lines with respect to each other (Westheimer and Hauske 1975; Mckee and Westhe 1978).

However, perceptual learning has also been implicated in the acquisition of the sorts of high-level properties of interest here. While the psychological study of the acquisition of subordinate-level categories (what is termed in that discipline ‘perceptual expertise’) has morphed into its own relatively independent subfield (though the field of radiological perceptual expertise is an exception, e.g. Sowden, Davies, and Roling 2000), much of the study of learned categorization more generally has remained within the domain of perceptual learning (Goldstone, Landy, and Brunel 2011; Gureckis and Goldstone 2008; Goldstone and Steyvers 2001; Goldstone 1994).

Perceptual learning for perceptual expertise involves processes known as attentional weighting and stimulus imprinting (Goldstone 1998; Goldstone and Byrge 2015).⁶⁷ In attentional

⁶⁷ I do not discuss unitization and differentiation – two other processes of perceptual learning – for the sake of simplicity. I do not believe this omission will impact any of the arguments that follow.

weighting, those features relevant to a given category are weighted in attention more heavily, and those that are irrelevant are weighted less heavily. The weights should be understood as evidential weights, where this means that their weighting corresponds to how diagnostic a given feature is for category membership. For example, eyes and face shape are highly diagnostic of category membership for barn owls, whereas the legs, colour and beak are diagnostic of category membership for blue herons. The processes by which this weighting is established is what is at issue here, and will be discussed in section 5.3.2. However, once the weighting is stable, attention to these features occurs in a covert and exogenous manner, and endogenous attention is not required.

Stimulus imprinting involves grouping the different attentionally weighted features together, as well as becoming sensitive to the relations between these features when they are relevant for diagnosticity. For example, what matters to categorizing faces as faces is not just the detection of eyes, a nose or a mouth, but detecting these features together and also in a specific configuration. The end result is a network of feature and feature-relation detectors that are grouped together.

This picture will strike readers as being very similar to the proposed explanation of diachronic cognitive permeation offered earlier, which I labeled the belief-driven hypothesis. However, it is one thing to say that perceptual learning involves the attentional re-weighting of features over time, and another to say that this amounts to cognitive permeation. I will argue in what follows that the belief-driven hypothesis is false: the learning of beliefs is not necessary for perceptual learning to occur. *Au contraire*, in some cases it is perceptual learning that allows for learning the beliefs that drive concept acquisition, amplification, or revision. I instead defend the *pattern-driven* hypothesis: perceptual learning is driven by real patterns in the relevant class of

training stimuli, and it is the gradual learning and detection of these patterns that drives the restructuring of the perceptual system, rather than declarative learning. Background beliefs are not necessary for perceptual learning to occur, and when cognition does guide learning, it does not count as cognitive permeation.

5.3.2 Arguing for the pattern-driven hypothesis

Some initial support for the pattern-driven hypothesis comes from the large numbers of exemplars needed for perceptual learning to occur. In the case of perceptual expertise, learning to categorize lab-manufactured objects such as greebles (gremlin-like animals) into subordinate categories requires roughly nine hours of practice with exemplars, broken into one-hour sessions spread out over two weeks (Gauthier and Tarr 1997; Gauthier et al. 1998). In well-studied domains of expertise such as chess, musical performance, and classical composing, the time required to develop ‘full blown’ expertise (where this likely outstrips perceptual expertise) is often estimated to be roughly a decade (Ericsson and Crutcher 1990; Hayes 1989; Weisberg 1999).

This exposure to exemplars is not optional to developing perceptual expertise. It suggests that the primary factor in assigning attentional weights is not what background beliefs one has, or one’s background theory, but rather what the statistical properties of the environment are – if a white stripe is important to categorizing something as a skunk, then given enough examples this feature will be weighted more heavily. This provides some support for the pattern-driven hypothesis over the belief-driven hypothesis, but it does not yet rule out possibility that beliefs play an essential role in perceptual learning. The belief-driven hypothesis may also be able to

explain the gradual nature of the process – while the background beliefs can cause shifts to our attention immediately, the perceptual system requires time to automate and synchronize the process (so that multiple features are selected simultaneously), requiring a certain number of pairings between the features. We need further reasons to favour one hypothesis over the other.

On the belief-driven hypothesis, acquisition of the relevant belief(s) is not a sufficient condition, but it is a necessary condition for perceptual expertise to occur. In order to assess this more easily, we can adapt Siegel's counterfactual analysis of cognitive permeation for the diachronic case (my changes in italics):

If the subject had not *learned the background belief(s) B* but was seeing and attending to the same distal stimuli, she would not have an experience with content p.

Attention here should be understood as spatial attention, rather than the feature-based attention involved in perceptual learning. In the latter instance, we cannot hold attention fixed between novice and expert in part because synchronous attention to the features relevant for categorization is itself a component of perceptual expertise. In the former case we can. Seeing a distal stimuli here should be understood in terms of low-level properties only, given that what is at issue is whether the expert can come to perceive high-level properties. Again, the counterfactual will not hold in all cases because there may be a closely related set of background beliefs B* that would allow the subject to also have an experience with content p. However, this analysis is useful because it makes clear that if the attentional reweighting involved in perceptual expertise – posited by both the pattern-driven and the belief-driven hypotheses – can occur

without being driven by the relevant background belief(s), then it is not an instance of cognitive permeation.

Attentional weighting absent belief can indeed occur. Goldstone, Landy, and Brunel (2011) characterize this sort of process as ‘blind flailing:’ attentional weights to features are assigned randomly, and those that allow the agent to make important or useful discriminations are selected and made permanent. This process is much like natural selection, where random variation plus a mechanism of selection can lead to strategic adaptations, without anyone having the intent of designing a creature with such an adaptation. The process is ‘blind’ in that there is no initial belief on the part of the agent that guides the process.

This sort of process has been well-studied in low-level perceptual learning, where such learning can occur absent cognitive guidance, in the sense that the learning occurs despite it being task-irrelevant – and so not guided by endogenous attention – or even when the stimuli are not consciously perceived (Watanabe, Náñez, and Sasaki 2001; Seitz and Watanabe 2009). However, it has been less well studied for the perceptual learning involved in perceptual expertise, as category learning is usually part of the assigned lab task.

One notable exception is (Y. K. Wong, Folstein, and Gauthier 2011), where subjects acquired perceptual expertise with respect to a lab-created object category called ‘ziggerins’ despite its irrelevance to the experimental task. Subjects were first given a periphery matching pre-test that involved being shown a target ziggerin presented in the center of the screen, followed by two ziggerins (one identical to the target, the other different in shape) presented in opposite diagonal corners of the screen while subjects were instructed to maintain fixation on a

cross in the center of the screen.⁶⁸ Subjects had to press a button to indicate the location of the target ziggerin. Subjects were then split into two groups and trained to either perform a task typical of perceptual expertise or perceptual learning experiments. For our purposes only the perceptual learning condition is relevant.

In the perceptual learning condition subjects were trained to perform a visual search task. First they were presented for 500 ms with an individual ziggerin rotated to the target angle. This was followed by the presentation of a circular array of eight ziggerins identical in shape to the initial ziggerin but rotated at different angles. Subjects then had to identify whether any of the ziggerins was rotated at the target angle. This was repeated with differently shaped ziggerins over a period of eight one-hour sessions. After the training period, a post-test identical to the pre-test was re-administered. The ability of the subjects to perform the post-test was significantly better after the perceptual learning training.

The experimenters take this as evidence for what they term ‘task-irrelevant perceptual expertise’ – perceptual expertise developed in the absence of training that requires its development. In each training sequence, all the individual ziggerin images were identical except for rotation. The subtle variations in shape between individual ziggerins were therefore irrelevant to the task. Nevertheless, subjects were passively exposed to ziggerins that varied in shape over the course of their training. The post-test was meant to test shape discrimination performance. This is a common marker of perceptual expertise with subordinate categories: the ability to make fine discriminations between similarly shaped individuals of the same basic-level category. The significant improvement of subjects in the post-test peripheral matching task – which involves

⁶⁸ Two other pre- and post-tests were given, but I exclude them from the discussion for simplicity.

the ability to make such fine discriminations – is therefore taken as evidence for task irrelevant perceptual expertise.

In addition, the sort of background belief that is posited by the belief-driven hypothesis is not present in this case. Experimenters did not provide any prior instruction to subjects as to what intrinsic features of the ziggerins were relevant to distinguishing between them, as only rotation was relevant to completing the task. So in this case no background beliefs or endogenous attention to category-distinguishing features are necessary for perceptual expertise to occur.

While one counterexample is sufficient to falsify the claim that background belief is necessary for perceptual expertise to develop, we may wonder how often blind flailing actually occurs in practice. Much perceptual learning for expertise takes place with some sort of cognitive guidance, or the adoption of some sort of strategy on the learner's part. In light of this, then perhaps we ought to make distinctions between different types of perceptual expertise, and evaluate the epistemic consequences separately, given their different etiologies. Perhaps background belief is necessary for some varieties of perceptual expertise.

Indeed, there are studies that suggest that some cognitive guidance is required (McGugin et al. 2011; Scott et al. 2006; Tanaka, Curran, and Sheinberg 2005; A. C.-N. Wong, Palmeri, and Gauthier 2009; A. C.-N. Wong et al. 2009). For example, (Stokes 2020 IV.2) cites a study by (Tanaka, Curran, and Sheinberg 2005) as demonstrating that mere exposure to perceptual stimuli is insufficient to produce perceptual expertise. In the study, experimenters trained twenty-one bird novices after providing them with a pre-test sequential matching task to measure their ability to categorize owls and wading birds at the subordinate level. For instance, an image of an eastern screech owl might be followed either by a distinct image of the same subordinate category ('eastern screech owl') or by an image from another subordinate category ('barn owl'). After

this, all subjects received training typical of perceptual expertise studies over seven consecutive days. Ten of the subjects learned to categorize owls at the subordinate level ('eastern screech owl') and wading birds at the basic level ('wading bird'). For the other eleven, the training was reversed; they learned to categorize wading birds at the subordinate level ('green heron') and owls at the basic level ('owl').

After this training period, the subjects performed the same pre-training matching task, along with two new matching tasks of the same kind as the first but with different images. The first new matching task contained novel instances of the same subordinate bird categories used in training; for example, a picture of a screech owl would be one not previously used in training. The second new matching task contained novel instances of new subordinate categories. For example, subjects would be shown pictures of the northern hawk owl, a subordinate category not previously learned in training.

In this post-test, subjects who had been trained at the subordinate level for wading birds did better at categorizing wading birds into subordinate level in all three tasks, with the same finding for subjects trained at the subordinate level for owls. However, what is really of interest is whether they did better at making subordinate-level categorizations for the category on which they did not receive subordinate-level training. Here, subjects did make some significant improvement in the first post-test, though it was below the improvements seen for those who received subordinate level training. In the second and third post-tests, there is no pre-training baseline to compare the results to – only the results of the other group. In both these post-tests, those untrained in the subordinate categories did worse than those with training, though the difference between the groups was smaller than in the case of the first pre-test. Tanaka, Curran, and Sheinberg discuss their results as follows:

These results highlight an important distinction between simple perceptual exposure and perceptual experience. In this study, participants were exposed to owl and wading birds and equal number of times. Yet their cognitive experience of those perceptual events was profoundly influenced by the category level used in the training task (Schyns, 1998).

(Tanaka, Curran, and Sheinberg 2005, 149–50)

Radiographers – or X-ray technicians – are another example provided by Stokes of cases where passive exposure does not seem to result in perceptual learning (he cites Nodine et al. 1999). Radiologists are physicians trained to interpret medical images. Their expertise is often thought to be partly perceptual in nature, due to the speed and accuracy with which they are able to make diagnoses from a single glance (Bilalic et al. 2016; Drew et al. 2013; B. S. Kelly et al. 2016; G. Wood et al. 2013; Kundel et al. 2007). Unlike radiologists, radiographers are not physicians and are not typically trained in how to diagnose medical images. Instead, their training consists in (amongst other elements such as anatomy) how to properly use the machines in order to get reliable images.⁶⁹ They therefore plausibly lack the background beliefs that would provide attentional guidance posited by the belief-driven hypothesis in order to form the perceptual category ‘tumour’ or ‘nodule,’ amongst other diagnostic categories. If they indeed lack perceptual expertise with these categories after a long period of exposure to images that contain them, then it looks like vindication that at least some forms of perceptual expertise require cognitive permeation.

⁶⁹ Training in medical image interpretation is available to radiographers, but this is usually additional to their degree or certificate requirements in radiography.

Manning et al. (2006) provide some support for this.⁷⁰ They did not find a statistically significant difference between the performance of eight novices at detecting pulmonary nodules and that of five experienced radiographers without training in chest image interpretation. This is in comparison to the superior performance of the same five radiographers after completing six months of training in chest image interpretation, whose accuracy was on par with the performance of eight experienced radiologists.

However, in deciding whether or not these two cases count against the pattern-driven hypothesis, we must be careful to properly characterize the nature of the cognitive influence. Not all guidance of perceptual learning by cognition will count as cognitive permeation. (Goldstone, Landy, and Brunel 2011, 8) call the strategic manipulation of the perceptual learning process ‘myopic flailing’ to characterize a process that is more efficient than random variation but that falls short of cognitive permeation. They provide a long list of ways in which we can ‘hack’ or tune our perceptual systems to accelerate perceptual learning. For example, we might put two images of an almost-identical viceroy and a monarch next to each other in order to compare them and spot the differences. Or, in wine tasting we might move the wine around our mouths to get a fuller and longer taste, or suck in some air to enhance its olfactory profile.

Adopting such strategies does not count as cognitive permeation first because these plausibly serve merely as catalysts, ensuring that our perceptual equipment gets the right input,

⁷⁰ It is less clear that (Nodine et al. 1999) - the study that Stokes cites – serves as support. There is no pure novice contrast group to expert radiologists, only radiographers (x-ray technicians) with 1-9 years imaging experience and 2nd and 3rd year medical residents with some training in reading images. Furthermore, the overall performance accuracy of radiographers was above chance and on a par with the performance of the medical residents (though below the performance of radiologists) (p.579). While, as Stokes notes, the radiographers performed below chance on one measure, the measure in question was calculated in terms of the fraction of true positives subtracted by the fraction of false positives, which was then plotted as a function of decision time. In this case, radiographers detected more false positives than true positives when they took relatively longer to come to a decision (p. 582, figure 5). This is in contrast to when they took less time to come to a decision. In such cases, 32% more of their decisions were true positives in comparison to residents (p. 580). That is, they outperformed residents (and radiologists) when making quick decisions. Finally, false positives were calculated by aggregating two different kinds of false positives: detecting a lesion where there was none, and categorizing a benign lesion as malignant. That radiographers perform worse on this task is perhaps not surprising, given that their job consists in making sure that the x-ray is of good quality, and part of this consists in making sure that any potentially suspicious areas are clearly imaged. That is, it may be part of their job to err on the side of false positives.

or enough practice with the right exemplars. Catalysts serve to speed a process that would have occurred anyways given enough time and the right conditions, so in such cases the guiding beliefs are not necessary for learning – it is not the case that if the subject had not learned background belief(s) B then she would not have perceptual experience with content p. In such cases, the subject could still come to have perceptual experience with content p, albeit after a longer time period.

Second, the guiding beliefs in this sort of myopic flailing do not observe the semantic coherence condition. We may adopt a given strategy in the belief that it will be helpful in learning to perceptually distinguish a given category of objects without possessing a background belief that the category is distinguished in virtue of certain characteristics, or even a belief as to which category is which. In such cases, there is a semantically arbitrary relationship between the belief and the resulting perceptual content. For example, I may put a picture of a viceroy and monarch side by side to facilitate perceptual learning without knowing in advance what the perceptual difference is between them, or which one is the viceroy and which the monarch (the training sets might be simply labeled ‘x’ and ‘y’). As it happens, it turns out that monarchs are distinguishable partly in virtue of a black stripe that cuts the bottom of the viceroy’s wings. But using the same strategy I might have learned to distinguish them in another way if they had different distinguishing features.

The case of (Tanaka, Curran, and Sheinberg 2005) may be explained in this way. First, the results are suggestive of the catalyst interpretation. While greater post-test gains were seen for the subordinate categories that subjects had been trained in (subjects trained to distinguish various wading birds were better at performing the same/different task for various wading birds), there was nevertheless significant improvement for the subordinate categories that the subjects

had only received basic-level training for. Subjects trained to distinguish various wading birds were better at performing the same/different task for various owl species that they had been previously exposed to during their training (Tanaka, Curran, and Sheinberg 2005, 148–49). This improvement suggests that there is some level of perceptual expertise that can be acquired with passive exposure, or task-irrelevance. While the degree of expertise is lower when compared with those who received subordinate-level training, it may nevertheless be the case that subjects gain more perceptual expertise with more exposure. This interpretation is also plausible with respect to the way people form expertise in real life. For example, when I first moved to Vancouver, I was not able to identify northern flicker woodpeckers on sight, but I soon learned to do so as they are abundant in my neighbourhood. I developed perceptual expertise long before I looked up the name of the bird.

Second, the sort of training that was given to subjects does not satisfy the semantic coherence constraint. Again, subjects were not given explicit instruction on the defining or typical perceptual characteristics of different basic-level or subordinate categories. Instead, they were given mere semantic labels, and feedback in terms of correct or incorrect semantic labeling during the training process. Subjects had to learn on their own, in the face of the many exemplars of birds shown, which features were relevant for categorization.

Here one might be tempted to characterize the relevant background belief along the lines of ‘that is a blue heron’ and so claim that it does meet the requirement. After all, it is a belief that is virtually identical to the resulting content of perceptual experience. However, notice that the belief concerns a mere semantic label. The person might have learned ‘that is a blue-footed booby’ instead, and still come to develop a perceptual category for blue herons on this basis. No content concerning the relevant perceptual characteristics of blue herons is conveyed in the belief

involving the semantic label (except of course in this case that the heron is blue – though really it is more grey than blue).

Finally, the most difficult cases of perceptual learning will be those where instruction does take the form of providing an explicitly articulated rule along the lines of ‘category C is characterized by perceptual features x, y, z.’ However, in such cases the semantic coherence condition is still not satisfied in any robust way. Instead, semantic coherence is lucky in a different way: the background belief and the perceptual experience are semantically coherent only because the belief happens to be true. Siegel’s original counterfactual analysis will not help us here. Instead, we need to consider what would happen in cases where the background beliefs are false. If I believe that the difference between viceroys and monarchs is the specific shade of orange, I would spend a lot of time focusing on contrasting the colour of these butterflies. But I would not be successful on this basis at being able to perceptually distinguish the two, and so would not develop perceptual expertise. Instead, with enough exposure to reliable training sets (and so long as my attention to colour is not too much of a distraction from the other features) this belief would likely be undone through perceptual learning and I would instead come to form a different belief about the low-level properties in virtue of which they are distinguishable. The semantic content of the belief is therefore not what is ultimately doing the work – rather, it is the statistical properties of the perceptual exemplars one is exposed to.

This sort of case should lead us to consider that learning may take place relatively separately in both the perceptual system and cognition. While we may acquire a background belief prior to becoming perceptual experts, this does not mean that the belief is responsible for the perceptual expertise. Instead, in some cases the direction of dependence may be reversed, with the perceptual category anchoring the reference of the semantic label, or any other cognitive

information associated with the object. This is consistent with several proposed philosophical solutions to the symbol grounding problem (Harnad 1990). Our abilities to perceptually categorize objects undergirds and serve cognition in a fundamental way. This is not to say that cognition has no role to play. In the next section I provide a sketch of how cognition can help guide perceptual learning for expertise. This also illuminates the case of the lack of radiological expertise after mere exposure.

5.3.3 The role of cognition in perceptual learning

Despite perceptual learning's distinctness from cognitive permeation, our beliefs about categories are nevertheless useful to perceptual learning. Most bird identification books include a description of the typical range, habitat and behavioural characteristics of the bird. For example, the killdeer:

is a gifted actor, well known for its "broken wing" distraction display. When an intruder wanders too close to its nest, it is greeted by an adult, who cries piteously while dragging a wing and stumbling about as if injured. Most predators take the bait and follow, and once the Killdeer has lured the predator far away from its nest, it miraculously recovers from the injury and flies off with a loud call.

(Campbell et al. 2005, 65)

This sort of information is pretty useful in identifying a killdeer if you happen too close to its nest. In such cases, it can serve to confirm the identity of the bird and so allow you to (non-

perceptually) categorize the bird as a killdeer. Knowing its habitat is also important. You would not usually find a killdeer in a forest, as it is mostly a shorebird, occupying beaches, mudflats and the like. This information may allow us to engage in non-perceptual categorization, which can then facilitate perceptual learning because it serves to confirm in another way that this is indeed an instance of a given type – it provides a form of simple feedback in much the same way that providing a reward or punishment for (un) successful categorization helps speed perceptual learning. So this also helps birders to self-create reliable training sets.

The role of cognition in perceptual learning is best understood as accomplishing the following:

1. Creating a reliable training set: either others do this for us via semantic labeling ('here is a file of all the viceroys, and here is one of all the monarchs'), or we do this for ourselves by drawing on other knowledge we possess ('this orange and black butterfly must be a monarch because it is feeding on milkweed').

2. Indicating that it is a relevant/important category: assigning a semantic label to something suggests that it has importance or relevance in our culture, and so we should pay attention to that thing.

3. Providing performance feedback: semantic labeling also serves as a sort of reward feedback to perceptual learning. The semantic information confirms that we have (un)successfully categorized an exemplar, and so can guide perceptual learning in this way.

4. Helping to select the right pattern: In cases where a list of perceptual characteristics is provided, this can help us to select the right pattern from all of those different patterns in the world that we could pick out. The world is teeming with patterns, and different ways of categorizing things. Some of these ways will be particularly useful or useless to us, and so

pattern selection can be determined in part by our goals and the real-world (dis)advantages of making or failing to make certain categorizations. But rather than learning exclusively through such trial and error, a list of typical characteristics allows us to attend to those features of the object that are part of the pattern. This is particularly useful in domains where perceptually distinguishable patterns are likely to be complex, such as with artworks.

5. Having goals: our desire to learn about a specific category of object is also very important to learning. If you have a complete disinterest in tree identification, then despite living in a city full of trees you may never come to be able to distinguish one subordinate category from another. However, if the categorization is interesting to you (perhaps because it is useful in some way) then this will initiate the sort of myopic flailing discussed above.

A better understanding of the role of cognition is helpful to understanding why radiographers may not develop perceptual expertise with tumours and nodules, despite extensive exposure. In such cases, it is likely that the divergent interests and goals of radiologists and radiographers have led them to focus on developing perceptual expertise with different categories by selecting different patterns. Whereas the goal of radiologists is to detect and diagnose potential abnormalities such as lesions and tumours, the goal of radiographers is to make sure the x-rays are good quality medical images. Considerable training is required in order to discern whether an image is of good quality. While there is no exact list of characteristics for what makes an image of good quality, the image properties radiographers are attuned to include image sharpness, density (degree of blackening on the film), and contrast (difference in density between two adjacent structures) (Easton 2009). So radiographers are likely perceptual experts with respect to whether a given x-ray is of good quality, though no studies have been done to test for such perceptual expertise.

However, if it is the case that radiographers do have such perceptual expertise, then it would go towards explaining their lack of expertise with respect to tumours and other abnormalities – during their exposure to images, they are searching for different patterns. Furthermore, unlike radiologists, radiographers do not typically receive feedback on which x-rays contain such abnormalities, and so have not had access to reliable training sets.⁷¹

5.4 Categorization vs. concepts

Some readers at this point might wonder how ‘perceptual’ categorization could be anything but cognitive. After all, categorization seems to require the possession and exercise of a concept, and since concepts are the stuff of thought, then categorization is instead properly a cognitive operation. While concepts may not be analyzable as a set of beliefs regarding necessary and sufficient conditions for application, they nevertheless seem to fall outside the boundaries of perception.

This line of reasoning may be behind the inferential leap from learning to cognitive permeation sketched earlier in section two.

For example, Fiona Macpherson writes: “the state of one’s cognitive system is determined, in part, by which concepts one possesses. If the possession of a concept affected one’s perceptual experience then that state of concept possession would be one that cognitively penetrated one’s experience” (2012, 27 fn. 3).

⁷¹ A further source of support for understanding human expertise with radiological images as perceptual can be found in machine learning, where deep neural networks are becoming increasingly adept at categorizing images as containing benign or malignant tumours only on the basis of having received training with labeled images (Shen et al. 2019). A full discussion of this line of support would take us far afield of the main task here, however.

John Zeimbekis and Athanasios Raftopoulos: “Siegel’s (2006; 2010) thesis that visual experience is affected by concept possession and that the contents of perception include high-level properties such as kinds implies a form of cognitive penetration of perception[...]” (2015a, 6).

Dustin Stokes: “some have recently argued that one’s experiences change by virtue of the acquisition of new recognitional capacities. This involves a kind of diachronic cognitive penetration, where one’s concepts—understood as abilities—affect how one perceives the world” (2015, 75).

However, this line of reasoning is too quick. As Siegel points out:

Empirical theories of object recognition are supposed to explain the nature of each of these components (the memory, the input, and the matching), and the mechanisms that underlie them. Part of what’s at issue in the debate about the Rich Content View is whether visual experience is only ever an input to such processes of recognition or whether it can also be an output.

(Siegel 2010, 110)

Assuming categorization (or recognition) is by default cognitive precludes two alternate possibilities. First, it may be the case that the visual system operates with its own proprietary perceptual concepts that are distinct from (though perhaps connected in important ways to) the sorts of concepts involved in thought. For example, it may yet be the case that while some concepts are indeed stored in long-term memory, and so count as cognitive, others are stored within the perceptual system itself. On this view there may be more than one way to recognize

objects: directly via perceptually stored concepts and indirectly via cognitively stored concepts that require some inference on the part of the perceiver in order to apply them. Indeed, Fodor (1983, 134) considered that something like this might be the case, both holding the view that perception is encapsulated from cognition, and proposing that the outputs of the perceptual system might be conceptual insofar as they output Eleanor Rosch's basic-level categories such as 'dog', 'bird' and the like (see also (Mandelbaum 2018) for a more sustained defence of this view).⁷²

The second closely related possibility is that our perceptual categories (where some of these correspond to high-level properties) are distinct entities from concepts. On this view, the perceptual system, while in the business of categorization, does not employ concepts to do so. The nature of perceptual categories is sufficiently distinct from concepts so as to warrant a different name. Pylyshyn's (1984; 1999) view might be an example. He discusses the possibility that through perceptual learning the visual system can come to recognize quite complex patterns without relying on general memory – what he calls 'compiled transducers.' Such complex patterns might at least co-vary with high-level properties, such as word meanings, and might themselves be counted as high-level properties. He writes that "there is no reason why the visual system could not encode any property whose identification does not require accessing long-term memory, and in particular that does not require inference from general knowledge" (1999, 361).

Which of these two options is more plausible will largely turn on how demanding we make our notion of a concept – the more demanding, the less likely we are to call recognition

⁷² For more on perceptual concepts see (Matthen 2005).

conceptual.⁷³ While I will not attempt to settle the issue here, it is sufficient to note that – given these alternatives – it will not do to simply assume that (perceptual) categorization is driven by or necessarily involves cognitive processes. Moreover, there are some reasons to reject this claim. A full exposition of these reasons is beyond the scope of this chapter, but I argue in chapter two that the speed at which expert categorization can take place, as well as the automaticity and encapsulation of the process all count against the idea that such categorization is a post-perceptual process, and instead favour the hypothesis that it takes place within perception; these characteristics are hallmarks of perceptual processing.

The proposal here that there is a distinction between perceptual categories used to perceptually categorize objects (where these may or may not be concepts) and cognitive concepts sits well with dual theories of concepts (E. E. Smith and Medin 1981; S. L. Armstrong, Gleitman, and Gleitman 1983; Osherson and Smith 1981). While these theories do not necessarily adopt the claim that categorization happens in perception, they nevertheless hold that concepts come in two varieties: ‘family resemblance’ concepts – or prototypes – that do not have necessary and sufficient conditions for application, and ‘classical’ concepts that do. They recognize these two entities as different kinds of concepts, sometimes ascribing them different functional roles. For example, prototypes may be ascribed the role of identifying objects, or as being involved in less rigorous reasoning, and classical concepts may be ascribed the role of being involved in more rigorous reasoning.⁷⁴

⁷³ See (Deroy 2013) for an argument that what she calls ‘multi-modal representations’ – which she takes to be responsible for some purported cases of cognitive penetration – are distinct from concepts and therefore not genuine cases of cognitive penetration. This position might be taken as a variant of the second option here. However, I take it that she would deny that these representations count as high-level properties. (Mandelbaum 2018) offers a version of the first option, and provides an argument that the representations count as conceptual based on a principled notion of what it is to be a concept.

⁷⁴ Evidence that is given in favour of these dual theories of concepts may also be used to provide some support for the dual view I propose here. It usually takes the following form: one variety of concept cannot accommodate the variety of empirical results in the study of concepts. For a review of such evidence see for example (Pinker and Prince 1999).

5.5 Conclusion

In this chapter I have responded to a challenge to the wider scope thesis: even if our perceptual experiences can be enriched via learning, this ‘perceptual expertise’ is an instance of cognitive permeation, and so therefore either our post-perceptual beliefs are either unjustified or only mediately justified. Against this I have argued that perceptual expertise occurs via perceptual learning, where this involves the formation of a perceptual category that allows for the representation of a high-level properties in perceptual experience, and this learning occurs primary in response to the statistical properties of the environment. As such, it does not require the learning of background beliefs of the kind involved in cognitive permeation. Additionally, when cognition is involved in guiding the learning process, it does so in a way that cannot be labeled cognitive permeation.

Finally, I have proposed the sorts of concepts involved in representing high-level properties in perceptual experience are not cognitive. Either there are perceptual concepts, or they are non-conceptual perceptual categories. On both accounts, these structures operate as part of the perceptual system in a way that is relatively independent from cognition. As such, they are not instances of concepts permeating perception, and so not an instance of cognitive permeation.

This does not yet constitute a complete defence of the wider scope thesis. Such a defence would require two things that I have not done here: (1) a specification of what it is about perceptual experience (or its etiology) that provides immediate justification to perceptual beliefs, and (2) a demonstration that the sorts of rich contents acquired in perceptual expertise also

possess this characteristic (Siegel 2017; Brogaard and Gatzia 2018; Chomanski and Chudnoff 2018; Chudnoff 2018).

Nevertheless, the account I have supported in this chapter – the pattern-driven hypothesis – sets the stage for this work. According to the pattern-driven hypothesis, the restructuring of the perceptual system can occur without the guidance of beliefs that would amount to cognitive permeation (and so threaten immediate justification), and instead responds to genuine patterns in the environment. The world constrains which categories can be perceptually learned, and their usefulness to us constrains which categories are actually perceptually learned. The route by which the contents of perceptual experience become enriched seems to involve an increased sensitivity to the environment. This is as James and Eleanor Gibson characterized perceptual learning. As a result of perceptual learning, perception “is progressively in *greater* correspondence with stimulation, not less” (J. Gibson and Gibson 1955a).

However, that perceptual learning for expertise involves greater sensitivity to the incoming stimulus is not to say that it provides an infallible foundation for perceptual beliefs. There are still ways that perceptual expertise can go wrong (for further discussion of bias in perceptual learning, see the conclusion of this dissertation, section 6.5.1).

Chapter 6: Conclusion

In this dissertation I have defended the perceptual expertise thesis: through perceptual learning, experts come to represent high-level properties in perceptual experience imperceptible to novices.

6.1 Using empirical methods to resolve the content debate

In chapter two, I provided an overview of a promising means of advancing the debate on the contents of perception – the common mechanism argument. It involves finding a property or effect of the perceptual processing of low-level properties and testing to see whether high-level properties also exhibit the property or effect. Two advantages of this strategy are that it does not rely on philosophical intuitions about phenomenology, and that it begins from common ground between sparse and rich content theorists: both sides agree that the processing of low-level properties is perceptual. I have argued that this second advantage is what allows the strategy to avoid the criticism that empirical methodology cannot settle the content debate because such methodology itself depends on theoretical commitments.

While some of the criteria do fall prey to the converse fallacy in that the perceptual effect in question is not proprietary to the perceptual system, the strategy is nevertheless useful for two reasons. First because its strength lies in demonstrating that high-level properties exhibit several of the effects. Second, because some effects are proprietary to the perceptual system when it comes to specific tasks, and so the results of these experiments are informative.

I have also argued that while there are sound experimental techniques for overcoming

one version of the low-level confound issue, the common mechanism argument does not have the resources to address a more substantive philosophical version of the challenge to natural kind properties involving twin earth considerations. This highlights a limitation of the strategy, and suggests that more philosophical work ought to be done to address this challenge.

Finally, I have highlighted three issues that arise when interpreting the results of studies employed by the common mechanism argument. First, there may be an evidential asymmetry for a given effect, where if a high-level property exhibits a given effect this counts strongly in favour of its being represented in perception, but if it does not exhibit the effect then this counts only weakly against it being represented in perception (or vice versa). Second, we ought to generalize to the narrowest subclass of properties needed in order to explain the empirical results, in part because studies on perceptual expertise suggest that the contents of perception will alter with experience. Third, in the face of conflicting results we may nevertheless favour the positive findings (that a high-level property is represented in perception) because there may be alternate task-specific explanations as to why the high-level property does not exhibit the effect, compatible with its being represented in perceptual experience. Asymmetry may also be useful in interpreting conflicting results.

In summary, the common mechanism argument represents a powerful tool for adjudicating the content debate when multiple criteria are invoked with respect to a given high-level property, though care must be taken to understand the limitations of the strategy, and to interpret the results.

6.2 What is learned in perceptual learning

In chapter three I advanced my main argument for the perceptual expertise thesis: through perceptual learning we can come to represent high-level properties in perceptual experience. My argument involves three steps. First, I draw on the mechanics of perceptual learning to argue that the properties really are high-level. Attentional weighting and stimulus imprinting are two processes by which the perceptual learning of categories occurs. Stimulus imprinting groups together the features that are characteristic of a given category (along with the relations between those features), and attentional weighting assigns weights to these features in accordance with how diagnostic they are for this category. The result is a stable perceptual category that is capable of ascribing high-level properties to objects. Positing a separate post-perceptual categorization process is superfluous. Moreover, understanding what is ascribed in terms of a mere assemblage of low-level properties cannot account for the pattern that our perceptual categories detect.

Second, I appeal to a variant of what I term the common mechanism argument in order to argue that these properties are represented in perception. The common mechanism argument begins from common ground between rich and sparse content theorists by appealing to effects of the perceptual processing of low-level properties. If such effects can also be found in the perceptual processing of high-level properties, then this is evidence in favour of the claim that those high-level properties are also represented in perception. Perceptual learning provides additional a way of running a variant of this argument because it provides a contrast between the way in which novices and experts process the high-level properties of expertise. Experts process such properties more quickly and automatically than novices, and adaptation effects are exhibited for some such properties. This change in processing is best explained by positing that the properties come to be represented in perceptual experience.

Third, I argue that representing high-level properties in perception alters the phenomenal character of perceptual experience – the perceptual categories we possess determine in part what it is like to perceive the world. I appeal to categorical perception as supporting evidence.

Categorical perception is an observable psychological shift in our perceptual experience due either to learned or innate perceptual categories, where objects that are members of the same perceptual category appear more similar, and those of different categories appear less similar.

This is sufficient to demonstrate that high-level properties can alter phenomenal character.

6.3 Learning to see beauty: a novel account of aesthetic perception

In chapter four I proposed a novel view of how we come to perceive beauty that accounts for how at least some aesthetic properties can be both learned and perceptual. On Waltonian Perceptualism, perceptual learning for expertise explains how we come to represent new categories of objects in perception, the fluency hypothesis explains how these categories lead to perceptual experiences of beauty and ugliness (and perhaps aesthetic properties more generally), and an understanding of what it is to be a perceptual expert allows us to say when those experiences are veridical. This account has the virtue of being able to satisfy several desiderata for a theory of aesthetic perception. Most importantly, it reconciles the puzzle of how knowledge can be required for the development of aesthetic expertise without compromising a perceptual view.

6.4 The epistemology of perceptual learning

In chapter five I defended the wider scope thesis: perceptual learning widens the scope of immediately justified perceptual beliefs. I have responded to a challenge to the wider scope thesis: even if our perceptual experiences can be enriched via learning, this ‘perceptual expertise’ is an instance of cognitive permeation, and so therefore either our post-perceptual beliefs are either unjustified or only mediately justified. Against this I have argued that perceptual expertise occurs via perceptual learning, where this involves the formation of a perceptual category that allows for the representation of a high-level properties in perceptual experience, and this learning occurs primary in response to the statistical properties of the environment. As such, it does not require the learning of background beliefs of the kind involved in cognitive permeation. Additionally, when cognition is involved in guiding the learning process, it does so in a way that cannot be labeled cognitive permeation.

Finally, I have proposed the sorts of concepts involved in representing high-level properties in perceptual experience are not cognitive. Either there are perceptual concepts, or they are non-conceptual perceptual categories. On both accounts, these structures operate as part of the perceptual system in a way that is relatively independent from cognition. As such, they are not instances of concepts permeating perception, and so not an instance of cognitive permeation.

6.5 Directions for future research

I have two main projects that I wish to pursue as a result of writing this dissertation. They are issues that I had to put aside while writing, but that nevertheless tugged at the corners of my mind as I was working on my chapters. Sometimes, when it all got too much, I would write up a paragraph or two of thoughts on the issue. Post-dissertation, I hope to expand on these thoughts.

6.5.1 Bias in perceptual learning

The first major project that I wish to undertake post-dissertation is to identify the varieties of bias involved in perceptual learning for expertise, and to understand when these may have morally problematic consequences. Recent psychological research has highlighted the extent to which our judgments and behaviours are influenced by unconscious racial and gender biases – also known as ‘implicit biases’. The phenomenon is widespread, occurring even when we explicitly disavow racism and sexism, or when we ourselves are members of the relevant gender or minority (Raymond 2013).

The social consequences are real. For example, implicit bias has been implicated in how positively we assess the qualifications of a person when making hiring decisions (Segrest Purkiss et al. 2006). It also accounts in part for the continued disproportionate police shootings of unarmed African Americans (Mekawi and Bresin 2015; Spencer, Charbonneau, and Glaser 2016), the failure of law enforcement and juries to assign appropriate weight to the testimony from visible minorities (Wiener, Wiener, and Grisso 1989; Kang and Lane 2010), and for different treatment outcomes for patients of different racial backgrounds (A. R. Green et al. 2007; Hall et al. 2015; FitzGerald and Hurst 2017). Philosophical work on implicit bias has identified and debated moral issues pertaining to such bias, as well as the extent to which we should trust our judgments, given the likely presence of such biases (D. Kelly and Roedder 2008; Saul 2013; Brownstein and Saul 2016a; 2016b).

While much has been said to elucidate the nature of the cognitive mechanism that underlies our tendency to engage in stereotypical thinking (Leslie 2007; 2008; 2017), as well as

the cultural transmission and learning of such stereotypes (Gelman, Ware, and Kleinberg 2010; M. Rhodes, Leslie, and Tworek 2012) little attention has been paid to the perceptual foundations of some implicit biases (Siegel 2017).

Perceptual learning for expertise provides us with the means to better perceive the world - when it goes right, that is. When the learning process is biased in some way, what we perceive may itself make a contribution to prejudicial behaviours, undergirding and reinforcing racial and gender discrimination in a vicious feedback loop. My proposal to identify the varieties of bias involved in perceptual learning for expertise is thus important to understanding the basis of some of our prejudicial behaviours.

This research will build on my dissertation by making use of my positive account of perceptual expertise in order to diagnose ways in which such expertise can be distorted or lacking. A thorough understanding of how a process works as intended is helpful to understanding how and where the process can break down. Only then can we cogently discuss questions of moral blame or responsibility, and formulate and assess potential solutions.

The first main goal of my project is to clearly identify, describe and catalogue the varieties of bias in perceptual learning for expertise. This is foundational work that has not yet been undertaken in the field, and it is crucial both to ensuring that people do not ‘talk past’ each other and to identifying and addressing cases of bias that are likely morally problematic. The importance of this step is underscored by the difficulties that have plagued debates in the ethics of machine learning due to a lack of clarity over what is meant by bias (Courtland 2018; Lipton and Steinhardt 2018). This has hampered progress in educating the public about the issue, and also in crafting effective technological fixes to address bias.

My approach is multidisciplinary, engaging with research in perceptual psychology and psychophysics, philosophy, and computer/information science. The first step is to draw on an emerging body of literature in computer/information science and philosophy that attempts to make more fine-grained distinctions between the different forms of bias alluded to in discussions of the ethics of machine learning (Danks and John London 2017; Barocas, Hardt, and Narayanan 2018; B. Green and Hu 2018). While in its most basic form to be biased is just to deviate from some standard, the standard itself may be statistical, moral, or legal. These different standards can come apart. The products of statistically biased learning may be unbiased relative to a moral standard, and vice versa. Danks & London (2017) make several other useful distinctions between kinds of biases orthogonal to the moral/statistical division, such as distinguishing between bias that may arise from the way information is processed (algorithmic processing bias) and that which may arise due to a misapplication of what is learned (transfer context bias).

These distinctions can be productively applied to taxonomizing bias in perceptual learning for expertise: adapting an extant taxonomy is a useful basis for beginning to understand the varieties of bias in perceptual learning, and for being able to specify exactly how they are biased. This first step involves specifying exactly how they are applicable, and identifying cases of perceptual learning that fall into these categories. For example, an ostensible case of algorithmic processing bias is categorical perception, which involves the perceptual ‘warping’ of the world according to the perceptual categories we possess; physical differences between things that are of the same category are less noticeable, and differences between two objects that are of different categories are more noticeable (Goldstone and Hendrickson 2010).

However, there is also opportunity for research in perceptual psychology and psychophysics to alter and expand the machine-learning framework. The second step of my

project is therefore to bring research from these fields to bear on the concept of bias. For example, many perceptual learning studies employ psychophysical tools to separate out two variables that can affect subject responses: sensitivity and response bias (Lippa and Goldstone 2001; Jones et al. 2015). Sensitivity measures the subject's ability to perceive or perceptually distinguish some aspect of the stimulus. Response bias is thought to measure how the subject may vary her response strategy (usually in accordance with an experiment's reward or punishment schema) while what she perceives is hypothesized to remain constant. This does not fit well into the taxonomy above, and so further work should be done to specify the nature of this bias, along with others that arise in perceptual learning. The results of this step are also likely to inform work that attempts to define and taxonomize bias in machine learning, as it introduces additional kinds of bias that may be relevant to the field. This may be helpful in identifying new ways in which reliance on algorithmic decision-making systems unintentionally promotes systemic injustice.

The second overarching goal of this project is to identify morally problematic cases of biases that arise from perceptual learning. In some sense, this falls out of the first phase of my project - cataloguing different forms of bias in perceptual learning includes identifying those cases that are biased relative to a moral standard. Nevertheless, the work involved in identifying distinctly moral biases in perceptual learning involves additional steps to those outlined above. Centrally, it requires connecting research in perceptual learning to implicit bias.

Research on the cross-race effect illustrates how fruitful this connection can be. The cross-race effect occurs when we are unable to perceptually recognize or distinguish individuals from races with which we are unfamiliar. It is a case of perceptual bias due to lack of perceptual expertise, where one is able to perceive fine-grained differences between the facial features of

same-race faces, but not other-race faces (Tanaka, Kiefer, and Bukach 2004; Michel, Caldara, and Rossion 2006; Hancock and Rhodes 2008; though see S. G. Young and Hugenberg 2012). The ethical consequences emerge most dramatically in criminal contexts, when an eyewitness or the victim of a crime misidentifies the other-race perpetrator. Eyewitness misidentification is the greatest unique source of wrongful convictions in North America and Great Britain (S. M. Smith, Stinson, and Prosser 2004). In everyday contexts, this deficit may lead to effects such as university professors subtly favouring same-race students. Understanding the perceptual basis for this biased behaviour allows for a relatively simple fix: the cross-race effect disappears when a person develops perceptual expertise with cross-race faces (Sangrigoli et al. 2005; McGugin et al. 2011). This in turn may provide support for proactive racial integration policies.

In completing this step I will focus on identifying other cases where perceptual biases may feed into, reinforce, or otherwise relate to implicit bias. Promising areas include nascent research on how gender perception may influence our subsequent judgments of a person's competence and character (Lick and Johnson 2013), and on how shooter bias may be a product of response bias in perceptual learning (Correll et al. 2014).

Finally, I will turn to the rich philosophical discussion of the ethical issues pertaining to implicit bias and moral responsibility and blameworthiness, applying these issues to bias in perceptual learning (Isaacs 1997; Harman 2011; Graham 2014; Mason 2015). The main question I will focus on here concerns whether or not we can be held morally responsible or blameworthy for our biases, given that they are not under our direct control and that we are often not aware of their influence on our beliefs (Brownstein and Saul 2016b).

This project is well poised to make a practical contribution to current 'de-biasing' training efforts to reduce problematic bias in various professional groups such as doctors, police

officers and judges (Larrick 2004; B. W. Smith and Slack 2015). It also contributes to the epistemology of perception: identifying how learned perception goes wrong is relevant to determining when we ought to trust beliefs formed via our senses. Finally, understanding bias in human perceptual learning can inform the study of machine learning - both in how we understand bias in this field, and also in how we ought to interact with automated processes in order to produce favourable social outcomes. Understanding human bias is vital to deciding when and how machine learning ought to be used to supplement human judgment.

6.5.2 Perceptual learning and theory-ladenness

The second direction for further research is a contribution to the philosophy of science and to further theoretical understanding of perceptual learning. While perceptual learning may not count as cognitive permeation (see chapter five), there is a further question about whether or not it renders perception theory-laden. Theory-ladenness and cognitive permeation are not quite the same thing. A permeating belief may fail to count as a ‘theory,’ and perception might be argued to be theory-laden without any cognitive permeation occurring (Churchland 1988).

Theory-laden perception should be understood as the claim that one’s explicitly held theories alter the contents of perceptual experience.⁷⁵ It was originally presented as a worry for scientific practice (Hanson 1961; Kuhn 1962). If theory somehow changes what one perceives, then two scientists operating with different theories may have different perceptual experiences when making some scientific observation. This is problematic for the idea that perceptual

⁷⁵ In fact Kuhn’s discussion extends beyond theory-laden perception to the incommensurability of scientific terms in different paradigms. This issue need not concern us here, as we are interested only in effects on perceptual experience.

experience provides a theory-neutral ground for adjudicating between theories. In the worst-case scenario, a scientist with theory A takes the observation to support her theory, while a scientist with theory B takes the observation to support her theory. Perceptual experience is no longer a means of adjudicating between theory A and theory B.

We might hold that perceptual learning for expertise counts as theory-laden if perceptual categories are in some way to the product of explicitly held scientific theories.⁷⁶ The role of theory must be substantive enough to assure the theory-ladenness of perceptual experience itself, rather than some post-perceptual process, but nevertheless fall short of cognitive permeation. Goldstone & Byrge (2015) argue that perceptual learning is theory-laden in just such a sense. They write that:

habitually executed tasks cause not only particular sensory features to be selectively attended to, but also feature creation (Schyns et al., 1998) [...] Perceptual learning that requires allocating attention to previously created features can count as theory-driven perceptual learning given the importance of goals and tasks on the original construction of perceptual features. [...] Although this diachronic change does not count as direct cognitive penetrability of perception, it entails that people with different experiences could have fundamentally different perceptual systems. This, in turn, is problematic for claims that equivalent training can equalize perceptual differences among scientists, and consequently problematic for the existence of a theory-neutral perceptual ground.

(Goldstone and Byrge 2015, 814)

⁷⁶ While there is a second possibility that perceptual categories as I have described them here themselves count as tacit theories, this is not the sort of theory-ladenness that is usually taken to be problematic for scientific practice, and so I will not consider it here.

The feature creation that they allude to in perceptual learning involves two functional mechanisms of perceptual learning: unitization and differentiation (see also Goldstone 1998). In unitization, perceptual features or properties are grouped together so that they become processed in perception as one unit.⁷⁷ This is most evident in the case of face perception, where it is more difficult to identify isolated features (such as a mouth or a nose) as belonging to an individual than it is to identify those features in the context of a face (Tanaka and Farah 1993). It is also more difficult to ignore part of the face in tasks (Gauthier and Bukach 2007). There is also evidence for the unitization of non-face objects of expertise (Gauthier, Tarr, and Bub 2010). Differentiation is the perceptual de-grouping of features or properties that were previously processed as units. The classic example of this is saturation and brightness, which are typically psychologically fused in our colour perceptions. However, it is possible to differentiate these two with training (B. Burns and Shepp 1988; Goldstone and Steyvers 2001).

Both differentiation and unitization change the way experts can allocate attention by grouping and ungrouping features or properties of objects. This can alter which perceptual categories are created, because it can lead to different properties being selected as relevant to categorization. The resulting picture is one where two scientists operating under different theories may have differing perceptual categories, and when confronted with the same perceptual observation will make different perceptual categorizations, and so come to represent different high-level properties in perceptual experience. Because such scientists have unitized or differentiated different features of an object, then even if they are shown the same training data they will parse it differently (A. C.-N. Wong et al. 2009).

⁷⁷ Unitization is very similar to stimulus imprinting, except that in unitization the features that are unitized must have been previously separate, whereas in stimulus imprinting this need not be the case.

While the worry for scientific observation is intuitively troubling from an epistemological standpoint, it is less clear how to think about the issue in terms of justification (Lyons 2011). I will not attempt to spell out the worry in terms of mediate or lack of justification here.⁷⁸ Nevertheless, one may worry that if perceptual learning leads to incommensurable scientific observations, that negative epistemic consequences follow. As Zeimbekis and Raftopoulos point out:

One would expect perceptual learning, which is a honing or improvement of perceptual discernment in a specific domain, to be a benign influence on perception and not to lead to incommensurability of the learner's percepts with those of other perceivers. Yet Kuhn argued that theories are incommensurable precisely on the grounds that observers presented with the same stimuli can see different things if their brain circuits are shaped differently by perceptual learning.

(Zeimbekis and Raftopoulos 2015a, 19)

Zeimbekis and Raftopoulos go on to suggest that one might argue against Kuhn that since perceptual learning is driven by the statistical properties of the environment rather than declarative learning, “if two subjects developed perceptual skills under different Kuhnian paradigms but were trained with the same data, they would end up seeing the same things. In that case, perceptual learning would not entail the epistemological consequences usually expected of theory-ladenness” (2015a, p19) (see also Raftopoulos 2001b).

⁷⁸ To elaborate on the view that post-perceptual beliefs formed as the result of theory laden perception are unjustified one might draw on the arguments of Robert Cummins (1998), who argues that philosophical intuitions that are the result of explicitly held theories lack any epistemic force. To argue they are only mediately justified, one might draw on claims such as John McDowell's (1982): “Theory can partly ground a claim to knowledge even in cases where it is not consciously brought to bear; as with a scientist who (as we naturally say) learns to see the movements of particles in some apparatus” (478).

However, as Goldstone & Byrge note in the quote above, such concordance is not guaranteed. Perceptual learning, insofar as it is guided by theoretical goals and tasks, may lead to the unitization and creation of different features and properties, even when trained on the same data. The problem is that there are too many different patterns that one might come to be able to perceptually recognize, and so cognitive guidance helps us to pick the pattern that will be useful in our context. If two scientists are operating in different cultural contexts, then they may end up becoming experts with respect to different patterns.

In this project my direction of initial investigation is to disarm the problems raised by theory-ladenness by distinguishing between perceptual learning resulting in our forming differing perceptual categories vs. it resulting in our forming incommensurable perceptual categories. Only the latter, I will argue, would result in problems for scientific theorizing, but this is not possible under the pattern-driven hypothesis.

The mere fact of differing observations is not epistemically worrisome, provided these observations do not conflict. This is because the observations, though different, are nevertheless compatible. First the patterns may be overlapping, in that they categorize objects at different levels of specificity. For example, an expert at identifying owls may make the observation that there is an owl in the tree. We may contrast this with someone who lacks perceptual expertise with owls but possesses it with respect to birds. This person may make a different observation that there is a bird in the tree. These two observations are different, and are the result of training guided by different goals (an interest in owls vs. an interest only in the category of birds), but

they are nevertheless compatible. This sort of case is ubiquitous – real-world perceptual expertise can be achieved in subordinate or basic-level categories, guided by one’s interests.⁷⁹

Second, different learned patterns may both be ways of achieving the same aim, in that they independently lead scientists to make the same categorization. For example, recent findings in machine learning suggest that some neural networks categorize objects largely on the basis of visual texture, rather than shape, as humans do (Geirhos et al. 2018). Their relative success at performing categorization tasks reveals that just as an object’s shape is relatively good guide to its category, so too is its texture. For example, while many species of animals have fur, there are enough differences between the specific texture of that fur in different species such that it can be a reliable guide to animal category. So these neural networks are plausibly relying on a different pattern to make the same perceptual categorization as we are.⁸⁰ Again, this is not troubling from an epistemic point of view – there is more than one way of identifying something over time.

A more problematic case would be if scientists were to make incommensurable observations as the result of perceptual learning. Consider the duck-rabbit illusion, where a drawing is ambiguous between a duck and a rabbit. If we suppose a real-life analogy, an animal is either a duck or a rabbit, not both. If one scientist were to observe a duck, and the other a rabbit, then their observations would be incommensurable. Supposing that our ability to represent objects as rabbits or ducks in perceptual experience is the result of perceptual learning, then this would be a problematic case were it to occur.

⁷⁹ An interesting real-world case of this sort may be: fieldwork by (Berlin 1992) finds that the Tzeltal Maya of Southern Mexico and the Aguaruna Jivaro of North Central Peru – two cultures that rely heavily on the land for sustenance – have subordinate-level expertise with the plants and trees of their respective regions, whereas the work of (Rosch et al., 1976) demonstrates that subjects living in the urban environment of Berkeley California tend to possess only basic-level expertise with trees.

⁸⁰ However, human vs. neural network performance at object categorization starts to diverge greatly when noise is introduced into the images because this affects texture much more than shape. One might wonder on this basis if they are making different categorizations, or just the same categorizations with different skill levels.

My goal in this project is therefore to develop an account of why this cannot occur (or, perhaps, overturn my initial assumptions and demonstrate how it can occur, and so vindicate worries of theory-ladenness). One way of going about this is to develop a more robust account of patterns, drawing on the work of (Andersen 2017; Dennett 1991). A more robust metaphysics of patterns may be able to help determine when such patterns are truly incommensurable.

References

- Allison, T., Aina Puce, Dennis D. Spencer, and Gregory McCarthy. 1999. "Electrophysiological Studies of Human Face Perception. I: Potentials Generated in Occipitotemporal Cortex by Face and Non-Face Stimuli." *Cerebral Cortex* 9 (5): 415–30.
- Alston, William P. 1983. "What's Wrong with Immediate Knowledge?" *Synthese* 55 (1): 73–95.
- Andersen, H. K. 2017. "Patterns, Information, and Causation." *The Journal of Philosophy* 114 (11): 592–622.
- Arcaro, Michael J, Peter F Schade, Justin L Vincent, Carlos R Ponce, and Margaret S Livingstone. 2017. "Seeing Faces Is Necessary for Face-Domain Formation." *Nature Neuroscience* 20 (10): 1404–12.
- Armstrong, K.M. 2011. "Covert Spatial Attention and Saccade Planning." In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, 78–96. New York: Oxford University Press.
- Armstrong, Sharon Lee, Lila R. Gleitman, and Henry Gleitman. 1983. "What Some Concepts Might Not Be." *Cognition* 13 (3): 263–308.
- Arnheim, Rudolf. 1954. *Art and Visual Perception*. Berkeley: University of California Press.
- Arstila, Valtteri. 2016. "Perceptual Learning Explains Two Candidates for Cognitive Penetration." *Erkenntnis* 81 (6): 1151–72.
- Avgerinou, Maria, and John Ericson. 1997. "A Review of the Concept of Visual Literacy." *British Journal of Educational Technology* 28 (4): 280–91.
- Aydede, Murat. 2019. "Is the Experience of Pain Transparent?" *Synthese* 196 (2): 677–708.
- Aydede, Murat, and Matthew Fulkerson. 2014. "Affect: Representationalists' Headache." *Philosophical Studies* 170 (2): 175–98.
- Ball, Karlene, and Robert Sekuler. 1982. "A Specific and Enduring Improvement in Visual Motion Discrimination." *Science (New York, N.Y.)* 218 (4573): 697–98.
- . 1987. "Direction-Specific Improvement in Motion Discrimination." *Vision Research* 27 (6): 953–65.
- Bargh, John A., and Paula Pietromonaco. 1982. "Automatic Information Processing and Social Perception: The Influence of Trait Information Presented Outside of Conscious Awareness on Impression Formation." *Journal of Personality and Social Psychology* 43 (3): 437–49.
- Barocas, Solon, Moritz Hardt, and Arvind Narayanan. 2018. *Fairness and Machine Learning*.

- Barrett, LF, and M Bar. 2009. "See It with Feeling: Affective Predictions during Object Perception." *Philosophical Transactions of the Royal Society B: Biological Sciences* 364 (1521): 1325–34.
- Barry, Ryan A, Katharine Graf Estes, and Susan M Rivera. 2015. "Domain General Learning: Infants Use Social and Non-Social Cues When Learning Object Statistics." *Frontiers in Psychology* 6 (May): 551.
- Bayne, Tim. 2009. "Perception and the Reach of Phenomenal Content." *The Philosophical Quarterly* 59 (236): 385–404.
- . 2016. "Gist!" *Proceedings of the Aristotelian Society* 116 (2): 107–26.
- Bender, John W. 1995. "General but Defeasible Reasons in Aesthetic Evaluation: The Particularist/Generalist Dispute." *The Journal of Aesthetics and Art Criticism* 53 (4): 379–92.
- Bentin, Shlomo, Truett Allison, Aina Puce, Erik Perez, and Gregory McCarthy. 1996. "Electrophysiological Studies of Face Perception in Humans." *Journal of Cognitive Neuroscience* 8 (6): 551–65.
- Bergqvist, Anna, and Robert Cowan. 2018. "Introduction." In *Evaluative Perception*, edited by Anna Bergqvist and Robert Cowan, 1–15. Oxford: Oxford University Press.
- Berlin, Brent. 1992. *Ethnobiological Classification: Principles of Categorization of Plants and Animals in Traditional Societies*. Princeton: Princeton University Press.
- Bilalic, Merim, Thomas Grottenhaler, Thomas Nagele, and Tobias Lindig. 2016. "The Faces in Radiological Images: Fusiform Face Area Supports Radiological Expertise." *Cerebral Cortex* 26 (3).
- Birkhoff, George. 1933. *Aesthetic Measure*. Cambridge, Massachusetts: Harvard University Press.
- Block, Ned. 2010. "Attention and Mental Paint." *Philosophical Issues* 20 (1): 23–63.
- . 2014. "Seeing-As in the Light of Vision Science." *Philosophy and Phenomenological Research* 89 (3): 560–72.
- . 2016. "Tweaking the Concepts of Perception and Cognition." *Behavioral and Brain Sciences* 39 (01): e232.
- Bodelón, Clara, Mazyar Fallah, and John H Reynolds. 2007. "Temporal Resolution for the Perception of Features and Conjunctions." *The Journal of Neuroscience* 27 (4): 725–30.

- Bonjour, Laurence. 1978. "Can Empirical Knowledge Have a Foundation?" *American Philosophical Quarterly* 15: 1–13.
- . 1985. *The Structure of Empirical Knowledge*. Cambridge, Massachusetts: Harvard University Press.
- Bornstein, M. H., and N. O. Korda. 1984. "Discrimination and Matching within and between Hues Measured by Reaction Times: Some Implications for Categorical Perception and Levels of Information Processing." *Psychological Research* 46 (3): 207–22.
- Braun, Jochen, and Bela Julesz. 1998. "Withdrawing Attention at Little or No Cost: Detection and Discrimination Tasks." *Perception & Psychophysics* 60 (1): 1–23.
- Brewer, Bill. 1999. *Perception and Reason*. Oxford: Oxford University Press.
- Briscoe, Robert. 2015. "Cognitive Penetration and the Reach of Phenomenal Content." In *The Cognitive Penetrability of Perception*, edited by John Zeimbekis and Athanassios Raftopoulos, 174–203. Oxford University Press.
- Brogaard, Berit. 2013. "Do We Perceive Natural Kind Properties?" *Philosophical Studies* 162 (1): 35–42.
- . 2018. "In Defense of Hearing Meanings." *Synthese* 195 (7): 2967–83.
- Brogaard, Berit, and Bartek Chomanski. 2015. "Cognitive Penetrability and High-Level Properties in Perception: Unrelated Phenomena?" *Pacific Philosophical Quarterly* 96 (4): 469–86.
- Brogaard, Berit, and Dimitria Electra Gatzia. 2018. "The Real Epistemic Significance of Perceptual Learning." *Inquiry* 61 (5–6): 543–58.
- Brown, Jonathon D., Natalie J. Novick, Kelley A. Lord, and Jane M. Richards. 1992. "When Gulliver Travels: Social Context, Psychological Closeness, and Self-Appraisals." *Journal of Personality and Social Psychology* 62 (5): 717–27.
- Brownstein, Michael, and Jennifer Saul, eds. 2016a. *Implicit Bias and Philosophy, Volume 1: Epistemology and Metaphysics*. Oxford University Press.
- . , eds. 2016b. *Implicit Bias and Philosophy, Volume 2: Moral Responsibility, Structural Injustice and Ethics*. Oxford University Press.
- Buchner, Axel. 1994. "Indirect Effects of Synthetic Grammar Learning in an Identification Task." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20 (3): 550–66.
- Burge, Tyler. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.

- . 2014. “Reply to Block: Adaptation and the Upper Border of Perception.” *Philosophy and Phenomenological Research* 89 (3): 573–83.
- Burns, Barbara, and Bryan E. Shepp. 1988. “Dimensional Interactions and the Structure of Psychological Space: The Representation of Hue, Saturation, and Brightness.” *Perception & Psychophysics* 43 (5): 494–507.
- Burns, Edward M., and W. Dixon Ward. 1978. “Categorical Perception—Phenomenon or Epiphenomenon: Evidence from Experiments in the Perception of Melodic Musical Intervals.” *The Journal of the Acoustical Society of America* 63 (2): 456–68.
- Butler, Andrea, Ipek Oruc, Christopher J. Fox, and Jason J.S. Barton. 2008. “Factors Contributing to the Adaptation Aftereffects of Facial Expression.” *Brain Research* 1191 (January): 116–26.
- Butterfill, Stephen. 2009. “Seeing Causing and Hearing Gestures.” *The Philosophical Quarterly* 59 (236): 405–28.
- Calder, Andrew J., Andrew W. Young, David I. Perrett, Nancy L. Etcoff, and Duncan Rowland. 1996. “Categorical Perception of Morphed Facial Expressions.” *Visual Cognition* 3 (2): 81–118.
- Campanella, S., A. Chrysochoos, and R. Bruyer. 2001. “Categorical Perception of Facial Gender Information: Behavioural Evidence and the Face-Space Metaphor.” *Visual Cognition* 8 (2): 237–62.
- Campbell, Wayne, Gregory Kennedy, Krista Kagume, and Carmen Adams. 2005. *Compact Guide to British Columbia Birds*. Vancouver: Lone Pine Publishing.
- Carrasco, Marisa. 2011. “Visual Attention: The Past 25 Years.” *Vision Research* 51 (13): 1484–1525.
- Carroll, N. 2002. “Aesthetic Experience Revisited.” *The British Journal of Aesthetics* 42 (2): 145–68.
- Carruthers, P. 2006. “Simple Heuristics Meet Massive Modularity.” *The Innate Mind: Culture and Cognition*.
- Cavedon-Taylor, Dan. 2017. “Reasoned and Unreasoned Judgement: On Inference, Acquaintance and Aesthetic Normativity.” *The British Journal of Aesthetics* 57 (1): 1–17.
- Cecchi, Ariel S. 2014. “Cognitive Penetration, Perceptual Learning and Neural Plasticity” 68 (1): 63–95.
- Cherry, E. Colin. 1953. “Some Experiments on the Recognition of Speech, with One and with Two Ears.” *The Journal of the Acoustical Society of America* 25 (5): 975–79.

- Chomanski, Bartek, and Elijah Chudnoff. 2018. "How Perception Generates, Preserves, and Mediates Justification." *Inquiry* 61 (5–6): 559–68.
- Chudnoff, Elijah. 2018. "The Epistemic Significance of Perceptual Learning." *Inquiry* 61 (5–6): 520–42.
- Churchland, Paul M. 1979. *Scientific Realism and the Plasticity of Mind*. Cambridge, Massachusetts: Cambridge University Press.
- . 1988. "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor." *Philosophy of Science* 55 (2): 167–87.
- Clark, Austen. 2000. *A Theory of Sentience*. Oxford University Press.
- Clercq, Rafael De. 2002. "The Concept of an Aesthetic Property." *Journal of Aesthetics and Art Criticism* 60 (2): 167–76.
- Clifford, Colin W.G. 2002. "Perceptual Adaptation: Motion Parallels Orientation." *Trends in Cognitive Sciences* 6 (3): 136–43.
- Connolly, Kevin. 2014. "Perceptual Learning and the Contents of Perception." *Erkenntnis* 79 (6): 1407–18.
- . 2019. *Perceptual Learning: The Flexibility of the Senses*. Oxford: Oxford University Press.
- Correll, Joshua, Sean M. Hudson, Steffanie Guillermo, and Debbie S. Ma. 2014. "The Police Officer's Dilemma: A Decade of Research on Racial Bias in the Decision to Shoot." *Social and Personality Psychology Compass* 8 (5): 201–13.
- Cosmides, Leda. 1989. "The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Selection Task." *Cognition* 31 (3): 187–276.
- Cosmides, Leda, John Tooby, Laurence Fiddick, and Gregory A Bryant. 2005. "Detecting Cheaters." *TRENDS in Cognitive Sciences* 28 (11): 505.
- Courtland, Rachel. 2018. "Bias Detectives: The Researchers Striving to Make Algorithms Fair." *Nature* 558 (7710): 357–60.
- Crutchfield, Parker. 2012. "Representing High-Level Properties in Perceptual Experience." *Philosophical Psychology* 25 (2): 279–94.
- Cummins, Robert. 1998. "Reflection on Reflective Equilibrium: The Psychology of Intuition and Its Role in Philosophical Inquiry." In *Rethinking Intuition*, edited by Michael DePaul and William Ramsey, 113–28. New York: Rowman & Littlefield.
- Curby, Kim M., and Bruno Rossion. 2010. "Competition between Face and Nonface Domains of

- Expertise.” In *Perceptual Expertise: Bridging Brain and Behavior*, edited by Isabel Gauthier, Michael J. Tarr, and Daniel Bub, 245–70. Oxford: Oxford University Press.
- Currie, Gregory. 1989. *An Ontology of Art*. New York: Palgrave Macmillan.
- Danks, David, and Alex John London. 2017. “Algorithmic Bias in Autonomous Systems.” In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, 4691–97.
- Danto, Arthur C. 2001a. “The Pigeon within Us All: A Reply to Three Critics.” *Journal of Aesthetics and Art Criticism* 59 (1): 39–44.
- Danto, Arthur C. 2001b. “Seeing and Showing.” *Source: The Journal of Aesthetics and Art Criticism*. Vol. 59.
- Delaney, M, G Anzures, J S Herbert, P C Quinn, A M Slater, J W Tanaka, K Lee, and O Pascalis. 2011. “Perceptual Training Prevents the Emergence of the Other Race Effect during Infancy.” *PLoS ONE* 6 (5): 19858.
- Delorme, A, G Richard, and Michèle Fabre-Thorpe. 2000. “Ultra-Rapid Categorisation of Natural Scenes Does Not Rely on Colour Cues: A Study in Monkeys and Humans.” *Vision Research* 40 (16): 2187–2200.
- Dennett, Daniel C. 1991. “Real Patterns.” *The Journal of Philosophy* 88 (1): 27–51.
- Deroy, Ophelia. 2013. “Object-Sensitivity versus Cognitive Penetrability of Perception.” *Philosophical Studies* 162 (1): 87–107.
- Desimone, Robert. 1998. “Visual Attention Mediated by Biased Competition in Extrastriate Visual Cortex.” Edited by G. W. Humphreys, J. Duncan, and A. Treisman. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 353 (1373): 1245–55.
- Dorsch, F. 2013. “Non-Inferentialism about Justification—The Case of Aesthetic Judgements.” *The Philosophical Quarterly* 63 (253): 660–82.
- Dosher, B a, and Z L Lu. 1998. “Perceptual Learning Reflects External Noise Filtering and Internal Noise Reduction through Channel Reweighting.” *Proceedings of the National Academy of Sciences of the United States of America* 95 (23): 13988–93.
- Dretske, Fred I. 1981. *Knowledge and the Flow of Information*. MIT Press.
- . 2015. “Perception versus Conception: The Goldilocks Test.” In *The Cognitive Penetrability of Perception*, edited by John Zeimbekis and Athanassios Raftopoulos, 163–73. Oxford: Oxford University Press.
- Drew, Trafton, Karla Evans, Melissa L-H Võ, Francine L Jacobson, and Jeremy M Wolfe. 2013.

- “Informatics in Radiology: What Can You See in a Single Glance and How Might This Guide Visual Search in Medical Images?” *Radiographics : A Review Publication of the Radiological Society of North America, Inc* 33 (1): 263–74.
- Duncan, John. 1998. “Converging Levels of Analysis in the Cognitive Neuroscience of Visual Attention.” Edited by G. W. Humphreys, J. Duncan, and A. Treisman. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 353 (1373): 1307–17.
- Durgin, F H, and D R Proffitt. 1996. “Visual Learning in the Perception of Texture: Simple and Contingent Aftereffects of Texture Density.” *Spatial Vision* 9 (4): 423–74.
- Easton, Suzanne. 2009. *An Introduction to Radiography*. Churchill Livingstone/Elsevier.
- Egan, Frances. 2012. “Representationalism.” In *The Oxford Handbook of Philosophy and Cognitive Science.*, edited by Eric Margolis, Richard Samuels, and Stephen Stich, 250–72. Oxford: Oxford University Press.
- Ellamil, M., J. M. Susskind, and A. K. Anderson. 2008. “Examinations of Identity Invariance in Facial Expression Adaptation.” *Cognitive, Affective, & Behavioral Neuroscience* 8 (3): 273–81.
- Enns, James T., and Ronald A. Rensink. 1991. “Preattentive Recovery of Three-Dimensional Orientation from Line Drawings.” *Psychological Review* 98 (3): 335–51.
- Ericsson, K. Anders (Karl Anders), and R.J. Crutcher. 1990. “The Nature of Exceptional Performance.” In *Life-Span Development and Behavior*, edited by D.L Featherman and R.M. Lerner, 187–217. Hillsdale, NJ: Erlbaum.
- Etcoff, Nancy L., and John J. Magee. 1992. “Categorical Perception of Facial Expressions.” *Cognition* 44 (3): 227–40.
- Findlay, J.M., and I.D. Gilchrist. 2003. *Active Vision: The Psychology of Looking and Seeing*. Oxford: Oxford University Press.
- Fingerhut, Joerg, and Jesse Prinz. 2018. “Grounding Evaluative Concepts.” *Philosophical Transactions of the Royal Society B: Biological Sciences* 373 (1752): 20170142.
- Fiorentini, Adriana, and Nicoletta Berardi. 1980. “Perceptual Learning Specific for Orientation and Spatial Frequency.” *Nature* 287 (5777): 43–44.
- Firestone, Chaz, and Brian J Scholl. 2014. “‘Top-Down’ Effects Where None Should Be Found: The El Greco Fallacy in Perception Research.” *Psychological Science* 25 (1): 38–46.
- . 2016. “Cognition Does Not Affect Perception: Evaluating the Evidence for ‘Top-down’ Effects.” *Behavioral and Brain Sciences* 39 (E229).

- Fischer, Rico, and Franziska Plessow. 2015. "Efficient Multitasking: Parallel versus Serial Processing of Multiple Tasks." *Frontiers in Psychology* 6: 1366.
- Fish, William. 2013. "High-Level Properties and Visual Experience." *Philosophical Studies* 162 (1): 43–55.
- FitzGerald, Chloë, and Samia Hurst. 2017. "Implicit Bias in Healthcare Professionals: A Systematic Review." *BMC Medical Ethics* 18 (1): 19.
- Fodor, Jerry A. 1983. *The Modularity of Mind*. Cambridge, Massachusetts: MIT Press.
- . 1988. "A Reply to Churchland's 'Perceptual Plasticity and Theoretical Neutrality.'" *Philosophy of Science* 55 (2): 188–98.
- Fox, Christopher J., and Jason J.S. Barton. 2007. "What Is Adapted in Face Adaptation? The Neural Representations of Expression in the Human Visual System." *Brain Research* 1127 (January): 80–89.
- Francuz, Piotr, Iwo Zaniwski, Paweł Augustynowicz, Natalia Kopiś, and Tomasz Jankowski. 2018. "Eye Movement Correlates of Expertise in Visual Arts." *Frontiers in Human Neuroscience* 12 (March): 87.
- Fudge, Robert. 2005. "A Vindication of Strong Aesthetic Supervenience." *Philosophical Papers* 34 (2): 149–71.
- Fugate, Jennifer M B. 2013. "Categorical Perception for Emotional Faces." *Emotion Review* 5 (1): 84–89.
- Fürst, Martina. 2017. "On the Limits of the Method of Phenomenal Contrast." *Journal of the American Philosophical Association* 3 (02): 168–88.
- Gangestad, Steven W., Randy Thornhill, and Ronald A. Yeo. 1994. "Facial Attractiveness, Developmental Stability, and Fluctuating Asymmetry." *Ethology and Sociobiology* 15 (2): 73–85.
- Gatzia, Dimitria Electra, and Berit Brogaard. 2017. "Pre-Cueing, Perceptual Learning and Cognitive Penetration." *Frontiers in Psychology* 8 (May): 739.
- Gauthier, Isabel, and Cindy Bukach. 2007. "Should We Reject the Expertise Hypothesis?" *Cognition* 103 (2): 322–30.
- Gauthier, Isabel, Tim Curran, Kim M. Curby, and Daniel Collins. 2003. "Perceptual Interference Supports a Non-Modular Account of Face Processing." *Nature Neuroscience* 6 (4): 428–32.
- Gauthier, Isabel, and Charles A. Nelson. 2001. "The Development of Face Expertise." *Current Opinion in Neurobiology*.

- Gauthier, Isabel, P Skudlarski, J C Gore, and A W Anderson. 2000. "Expertise for Cars and Birds Recruits Brain Areas Involved in Face Recognition." *Nature Neuroscience* 3 (2): 191–97.
- Gauthier, Isabel, and Michael J. Tarr. 1997. "Becoming a 'Greeble' Expert: Exploring Mechanisms for Face Recognition." *Vision Research* 37 (12): 1673–82.
- Gauthier, Isabel, Michael J. Tarr, and Daniel Bub, eds. 2010. *Perceptual Expertise: Bridging Brain and Behavior*. Oxford: Oxford University Press.
- Gauthier, Isabel, Pepper Williams, Michael J. Tarr, and James Tanaka. 1998. "Training 'Greeble' Experts: A Framework for Studying Expert Object Recognition Processes." *Vision Research* 38 (15–16): 2401–28.
- Geirhos, Robert, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. 2018. "ImageNet-Trained CNNs Are Biased towards Texture; Increasing Shape Bias Improves Accuracy and Robustness."
- Gelman, Susan A., Elizabeth A. Ware, and Felicia Kleinberg. 2010. "Effects of Generic Language on Category Content and Structure." *Cognitive Psychology* 61 (3): 273–301.
- Gibson, Eleanor J. 1963. "Perceptual Learning." *Annual Review of Psychology* 14 (January): 29–56.
- Gibson, Eleanor J., and Anne D Pick. 2000. *An Ecological Approach to Perceptual Learning and Development*. Oxford University Press.
- Gibson, James, and Eleanor J. Gibson. 1955a. "Perceptual Learning: Differentiation or Enrichment?" *Psychological Review*.
- . 1955b. "What Is Learned in Perceptual Learning? A Reply to Professor Postman." *Psychological Review* 62 (6): 447–50.
- Gijp, A. van der, C. J. Ravesloot, H. Jarodzka, M. F. van der Schaaf, I. C. van der Schaaf, J. P. J. van Schaik, and Th. J. ten Cate. 2017. "How Visual Search Relates to Visual Diagnostic Performance: A Narrative Systematic Review of Eye-Tracking Research in Radiology." *Advances in Health Sciences Education* 22 (3): 765–87.
- Gilbert, Cd, Mariano Sigman, and Re Crist. 2001. "The Neural Basis of Perceptual Learning." *Neuron* 31: 681–97.
- Goldie, Peter., and Elisabeth. Schellekens. 2009. *Who's Afraid of Conceptual Art?* Routledge.
- Goldman, Alvin I. 2008. "Immediate Justification and Process Reliabilism." In *Epistemology: New Essays*, edited by Quentin Smith, 63–82. Oxford University Press.
- Goldstone, Robert L. 1994. "Influences of Categorization on Perceptual Discrimination."

- Journal of Experimental Psychology: General*.
- . 1998. “Perceptual Learning.” *Annual Review of Psychology*.
- Goldstone, Robert L., and LW Barsalou. 1998. “Reuniting Perception and Conception.” *Cognition*.
- Goldstone, Robert L., and Lisa A. Byrge. 2015. “Perceptual Learning.” In *The Oxford Handbook of Philosophy of Perception*, edited by Mohan Matthen, 812–32. Oxford: Oxford University Press.
- Goldstone, Robert L., and Andrew T. Hendrickson. 2010. “Categorical Perception.” *Wiley Interdisciplinary Reviews: Cognitive Science* 1 (1): 69–78.
- Goldstone, Robert L., David Landy, and Lionel C. Brunel. 2011. “Improving Perception to Make Distant Connections Closer.” *Frontiers in Psychology* 2 (December): 385.
- Goldstone, Robert L., and Mark Steyvers. 2001. “The Sensitization and Differentiation of Dimensions during Category Learning.” *Journal of Experimental Psychology: Learning, Memory, and Cognition* 130 (1): 116–39.
- Gombrich, Ernst Hans. 1980. *The Sense of Order*. London: Phaidon.
- Gordon, Peter C., and Keith J. Holyoak. 1983. “Implicit Learning and Generalization of the ‘Mere Exposure’ Effect.” *Journal of Personality and Social Psychology* 45 (3): 492–500.
- Graham, Peter A. 2014. “A Sketch of a Theory of Moral Blameworthiness*.” *Philosophy and Phenomenological Research* 88 (2): 388–409.
- Green, Alexander R., Dana R. Carney, Daniel J. Pallin, Long H. Ngo, Kristal L. Raymond, Lisa I. Iezzoni, and Mahzarin R. Banaji. 2007. “Implicit Bias among Physicians and Its Prediction of Thrombolysis Decisions for Black and White Patients.” *Journal of General Internal Medicine* 22 (9): 1231–38.
- Green, Ben, and Lily Hu. 2018. “The Myth in the Methodology: Towards a Recontextualization of Fairness in Machine Learning.” In *The ICML 2018 Debates Workshop*.
- Greene, Michelle R., and Oliva. 2009. “The Briefest of Glances: The Time Course of Natural Scene Understanding.” *Psychological Science* 20 (4): 464–72.
- Greene, Michelle R., and Aude Oliva. 2010. “High-Level Aftereffects to Global Scene Properties.” *Journal of Experimental Psychology: Human Perception and Performance* 36 (6): 1430–42.
- Gross, Steven. 2017. “Cognitive Penetration and Attention.” *Frontiers in Psychology* 8: 1–12.

- Gureckis, Todd M, and Robert L. Goldstone. 2008. "The Effect of the Internal Structure of Categories on Perception." In *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, 1876–81.
- Halberstadt, Jamin, and Gillian Rhodes. 2000. "The Attractiveness of Nonface Averages: Implications for an Evolutionary Explanation of the Attractiveness of Average Faces." *Psychological Science* 11 (4): 285–89.
- . 2003. "It's Not Just Average Faces That Are Attractive: Computer-Manipulated Averageness Makes Birds, Fish, and Automobiles Attractive." *Psychonomic Bulletin & Review* 10 (1): 149–56.
- Hall, William J, Mimi V Chapman, Kent M Lee, Yesenia M Merino, Tainayah W Thomas, B Keith Payne, Eugenia Eng, Steven H Day, and Tamera Coyne-Beasley. 2015. "Implicit Racial/Ethnic Bias Among Health Care Professionals and Its Influence on Health Care Outcomes: A Systematic Review." *American Journal of Public Health* 105 (12): e60-76.
- Hancock, Kirsten J., and Gillian Rhodes. 2008. "Contact, Configural Coding and the Other-Race Effect in Face Recognition." *British Journal of Psychology* 99 (1): 45–56.
- Hanson, Norwood R. 1961. *Patterns of Discovery*. Cambridge, Massachusetts: Cambridge University Press.
- Harman, Elizabeth. 2011. "Does Moral Ignorance Exculpate?" *Ratio* 24 (4): 443–68.
- Harnad, Stevan, ed. 1987. *Categorical Perception: The Groundwork of Cognition*. Cambridge: Cambridge University Press.
- . 1990. "The Symbol Grounding Problem." *Physica D: Nonlinear Phenomena* 42 (1–3): 335–46.
- . 2003. "Categorical Perception." In *Encyclopedia of Cognitive Science*. Nature Publishing Group.
- Hayes, J.R. 1989. *The Complete Problem Solver*. 2nd ed. Hillsdale, NJ: Erlbaum.
- Hekkert, P., and P. C. W. Wieringen. 1990. "Complexity and Prototypicality as Determinants of the Appraisal of Cubist Paintings." *British Journal of Psychology* 81 (4): 483–95.
- Helton, Grace. 2016. "Recent Issues in High-Level Perception." *Philosophy Compass* 11 (12): 851–62.
- Hershler, Orit, and Shaul Hochstein. 2005. "At First Sight: A High-Level Pop out Effect for Faces." *Vision Research* 45 (13): 1707–24.
- Hooks, Bryan M., and Chnfei Chen. 2007. "Critical Periods in the Visual System: Changing Views for a Model of Experience-Dependent Plasticity." *Neuron* 56 (2): 312–26.

- Hopkins, Robert. 2005. "Aesthetics, Experience, and Discrimination." *The Journal of Aesthetics and Art Criticism* 63 (2): 119–33.
- . 2006. "Painting, History, and Experience." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 127 (1): 19–35.
- Hubel, D. H., and T. N. Wiesel. 1970. "The Period of Susceptibility to the Physiological Effects of Unilateral Eye Closure in Kittens." *The Journal of Physiology* 206 (2): 419–36.
- Huemer, Michael. 2001. *Skepticism and the Veil of Perception*. Rowman & Littlefield Publishers.
- Humphrey, Diane. 1997. "Preferences in Symmetries and Asymmetries in Drawings: Asymmetries between Ages and Sexes." *Empirical Studies of the Arts* 15 (1): 41–60.
- Hutcheson, Francis. 2008. *An Inquiry into the Original of Our Ideas of Beauty and Virtue in Two Treatises*. Edited by Wolfgang Leidhold. Revised Ed. Indianapolis: Liberty Fund.
- Irvin, Sherri. 2008a. "The Pervasiveness of the Aesthetic in Ordinary Experience." *British Journal of Aesthetics* 48 (1): 29–44.
- . 2008b. "The Pervasiveness of the Aesthetic in Ordinary Experience." *The British Journal of Aesthetics* 48 (1): 29–44.
- Isaacs, Tracy. 1997. "Cultural Context and Moral Responsibility." *Ethics* 107 (4): 670–84.
- Iseminger, Gary. 2004. *The Aesthetic Function of Art*. Cornell University Press.
- Itz, Marlena L., Stefan R. Schweinberger, Claudia Schulz, and Jürgen M. Kaufmann. 2014. "Neural Correlates of Facilitations in Face Learning by Selective Caricaturing of Facial Shape or Reflectance." *NeuroImage* 102 (November): 736–47.
- Jolicoeur, Pierre, Mark A. Gluck, and Stephen M. Kosslyn. 1984. "Pictures and Names: Making the Connection." *Cognitive Psychology* 16 (2): 243–75.
- Jones, Pete R., David R. Moore, Daniel E. Shub, and Sygal Amitay. 2015. "The Role of Response Bias in Perceptual Learning." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 41 (5): 1456–70.
- Jonides, John, and Henry Gleitman. 1972. "A Conceptual Category Effect in Visual Search: O as Letter or as Digit." *Perception & Psychophysics* 12 (6): 457–60.
- Kang, Jerry, and Kristin Lane. 2010. "Seeing through Colorblindness: Implicit Bias and the Law." *UCLA Law Review* 58 (2): 465–520.
- Karni, A., and D. Sagi. 1991. "Where Practice Makes Perfect in Texture Discrimination: Evidence for Primary Visual Cortex Plasticity." *Proceedings of the National Academy of Sciences of*

- the United States of America* 88 (11): 4966–70.
- Kellman, Philip J., and Patrick Garrigan. 2009. “Perceptual Learning and Human Expertise.” *Physics of Life Reviews*.
- Kellman, Philip J, and Christine M Massey. 2013. “Perceptual Learning , Cognition , and Expertise.” In *Psychology of Learning and Motivation*, 58:117–65. Elsevier.
- Kelly, Brendan S., Louise A. Rainford, Sarah P. Darcy, Eoin C. Kavanagh, and Rachel J. Toomey. 2016. “The Development of Expertise in Radiology: In Chest Radiograph Interpretation, ‘Expert’ Search Pattern May Predate ‘Expert’ Levels of Diagnostic Accuracy for Pneumothorax Identification.” *Radiology* 280 (1): 252–60.
- Kelly, Daniel, and Erica Roedder. 2008. “Racial Cognition and the Ethics of Implicit Bias.” *Philosophy Compass* 3 (3): 522–40.
- Kirchner, Holle, and Simon J. Thorpe. 2006. “Ultra-Rapid Object Detection with Saccadic Eye Movements: Visual Processing Speed Revisited.” *Vision Research* 46 (11): 1762–76.
- Kleffner, Dorothy A., and V. S. Ramachandran. 1992. “On the Perception of Shape from Shading.” *Perception & Psychophysics* 52 (1): 18–36.
- Koksvik, Ole. 2015. “Phenomenal Contrast: A Critique.” *American Philosophical Quarterly* 52 (4): 321–34.
- Kompaniez-Dunigan, Elysse, Craig K. Abbey, John M. Boone, and Michael A. Webster. 2015. “Adaptation and Visual Search in Mammographic Images.” *Attention, Perception, & Psychophysics* 77 (4): 1081–87.
- Kompaniez, Elysse, Craig K. Abbey, John M. Boone, and Michael A. Webster. 2013. “Adaptation Aftereffects in the Perception of Radiological Images.” Edited by Samuel G. Solomon. *PLoS ONE* 8 (10): e76175.
- Kornblith, Hilary. 2002. *Knowledge and Its Place in Nature*. Oxford: Oxford University Press.
- Krueger, Lester E. 1984. “The Category Effect in Visual Search Depends on Physical Rather than Conceptual Differences.” *Perception & Psychophysics* 35 (6): 558–64.
- Kuhn, Thomas S. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Kundel, Harold L., Calvin F. Nodine, Emily F. Conant, and Susan P. Weinstein. 2007. “Holistic Component of Image Perception in Mammogram Interpretation: Gaze-Tracking Study.” *Radiology* 242 (2): 396–402.
- Kvanvig, Jonathan L., and Wayne D. Riggs. 1992. “Can a Coherence Theory Appeal to Appearance States?” *Philosophical Studies* 67 (3): 197–217.

- Lakoff, George, and Mark Johnson. 2008. *Metaphors We Live By*. Chicago: University of Chicago Press.
- Lamarque, Peter. 2010. *Work and Object: Explorations in the Metaphysics of Art*. Oxford: Oxford University Press.
- Lamme, Victor A.F., and Pieter R. Roelfsema. 2000. "The Distinct Modes of Vision Offered by Feedforward and Recurrent Processing." *Trends in Neurosciences* 23 (11): 571–79.
- Lane, Harlan. 1965. "The Motor Theory of Speech Perception: A Critical Review." *Psychological Review* 72 (4): 275–309.
- Langlois, Judith H., and Lori A. Roggman. 1990. "Attractive Faces Are Only Average." *Psychological Science* 1 (2): 115–21.
- Larrick, Richard P. 2004. "Debiasing." In *Blackwell Handbook of Judgment and Decision Making*, edited by Derek J. Koehler and Nigel Harvey, 316–38. John Wiley & Sons.
- Leopold, David A., Alice J. O'Toole, Thomas Vetter, and Volker Blanz. 2001. "Prototype-Referenced Shape Encoding Revealed by High-Level Aftereffects." *Nature Neuroscience* 4 (1): 89–94.
- Leslie, Sarah-Jane. 2007. "Generics and the Structure of the Mind." *Philosophical Perspectives* 21: 375–403.
- . 2008. "Generics: Cognition and Acquisition." *Philosophical Review* 117 (1): 1–47.
- . 2017. "The Original Sin of Cognition." *Journal of Philosophy* 114 (8): 393–421.
- Levin, Daniel T., and Bonnie L. Angelone. 2001. "Visual Search for a Socially Defined Feature: What Causes the Search Asymmetry Favoring Cross-Race Faces?" *Perception & Psychophysics* 63 (3): 423–35.
- Levin, Daniel T., and Bonnie L. Angelone. 2002. "Categorical Perception of Race." *Perception* 31 (5): 567–78.
- Levinson, Jerrold. 2001. "Aesthetic Properties, Evaluative Force and Differences of Sensibility." In *Aesthetic Concepts: Essays after Sibley*, edited by Emily Brady and Jerrold Levinson, 61–80. Oxford: Clarendon Press.
- . 2005. "Aesthetic Properties." *Proceedings of the Aristotelian Society, Supplementary Volumes* 78: 211–27.
- Li, Fei Fei, Rufin VanRullen, Christof Koch, and Pietro Perona. 2002. "Rapid Natural Scene Categorization in the near Absence of Attention." *Proceedings of the National Academy of Sciences of the United States of America* 99 (14): 9596–9601.

- Liberman, Alvin M., Katherine Safford Harris, Howard S. Hoffman, and Belver C. Griffith. 1957. "The Discrimination of Speech Sounds Within and Across Phoneme Boundaries." *Journal of Experimental Psychology* 54 (5): 358–68.
- Lick, David J., and Kerri L. Johnson. 2013. "Fluency of Visual Processing Explains Prejudiced Evaluations Following Categorization of Concealable Identities." *Journal of Experimental Social Psychology* 49 (3): 419–25.
- Lier, Jens Van, Russell Revlin, and Wim De Neys. 2013. "Detecting Cheaters without Thinking: Testing the Automaticity of the Cheater Detection Module." Edited by Tiziana Zalla. *PLoS ONE* 8 (1): e53827.
- Limbach, Katharina, Jürgen M. Kaufmann, Holger Wiese, Otto W. Witte, and Stefan R. Schweinberger. 2018. "Enhancement of Face-Sensitive ERPs in Older Adults Induced by Face Recognition Training." *Neuropsychologia* 119 (October): 197–213.
- Linkenkaer-Hansen, Klaus, Jaakko M Palva, Mikko Sams, Jari K Hietanen, Hannu J Aronen, and Risto J Ilmoniemi. 1998. "Face-Selective Processing in Human Extrastriate Cortex around 120 Ms after Stimulus Onset Revealed by Magneto- and Electroencephalography." *Neuroscience Letters* 253 (3): 147–50.
- Lippa, Yvonne, and Robert L. Goldstone. 2001. "The Acquisition of Automatic Response Biases through Categorization." *Memory & Cognition* 29 (7): 1051–60.
- Lipton, Zachary C., and Jacob Steinhardt. 2018. "Troubling Trends in Machine Learning Scholarship." In *ICML 2018: The Debates*.
- Liu, Hesheng, Yigal Agam, Joseph R. Madsen, and Gabriel Kreiman. 2009. "Timing, Timing, Timing: Fast Decoding of Object Information from Intracranial Field Potentials in Human Visual Cortex." *Neuron* 62 (2): 281–90.
- Liu, Jiajuan, Barbara Anne Doshier, and Zhong-Lin Lu. 2015. "Augmented Hebbian Reweighting Accounts for Accuracy and Induced Bias in Perceptual Learning with Reverse Feedback." *Journal of Vision* 15 (10).
- Livingston, Kenneth R., Janet K. Andrews, and Stevan Harnad. 1998. "Categorical Perception Effects Induced by Category Learning." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 24 (3): 732–53.
- Lloyd, David R. 2010. "Symmetry and Beauty in Plato." *Symmetry* 2 (2): 455–65.
- Logue, Heather. 2013. "Visual Experience of Natural Kind Properties: Is There Any Fact of the Matter?" *Philosophical Studies* 162 (1): 1–12.
- . 2018. "Can We Visually Experience Aesthetic Properties?" In *Evaluative Perception*, edited by Anna Bergqvist and Robert Cowan, 42–57. Oxford: Oxford University Press.

- Longworth, Guy. 2008. "Linguistic Understanding and Knowledge." *Noûs* 42 (1): 50–79.
- Lopes, Dominic McIver. 2014a. *Beyond Art*. Oxford University Press.
- . 2014b. "Feckless Reason." In *Aesthetics and the Sciences of Mind*, edited by G Currie, M Kieran, A Meskin, and J Robson, 21–36. New York: Oxford University Press.
- . 2016. "In the Eye of the Beholder." In *Art, Mind, and Narrative: Themes from the Work of Peter Goldie*, edited by Julian Dodd. Oxford: Oxford University Press.
- . 2018. *Being for Beauty: Aesthetic Agency and Value*. Oxford: Oxford University Press.
- Luca, Emanuela De, and Manfred Fahle. 1999. "Learning of Interpolation in 2 and 3 Dimensions." *Vision Research* 39 (12): 2051–62.
- Lycan, William G. 2014. "What Does Vision Represent?" In *Does Perception Have Content?*, edited by Berit Brogaard, 311–28. Oxford: Oxford University Press.
- Lyons, Jack. 2005. "Perceptual Belief and Nonexperiential Looks." *Philosophical Perspectives* 19 (1): 237–56.
- . 2011. "Circularity, Reliability, and the Cognitive Penetrability of Perception." *Philosophical Issues* 21: 289–311.
- Machery, Edouard. 2015. "Cognitive Penetrability: A No-Progress Report." In *The Cognitive Penetrability of Perception : New Philosophical Perspectives*, edited by John Zeimbekis and Athanassios Raftopoulos, 59–74.
- Macpherson, Fiona. 2012. "Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism." *Philosophy and Phenomenological Research* 84 (1): 24–62.
- Mandelbaum, Eric. 2015. "The Automatic and the Ballistic: Modularity beyond Perceptual Processes." *Philosophical Psychology* 28 (8): 1147–56.
- . 2018. "Seeing and Conceptualizing: Modularity and the Shallow Contents of Perception." *Philosophy and Phenomenological Research* 97 (2): 267–83.
- Manning, David, Susan Ethell, Tim Donovan, and Trevor Crawford. 2006. "How Do Radiologists Do It? The Influence of Experience and Training on Searching for Chest Nodules." *Radiography* 12 (2): 134–42.
- Manza, L., D. Zizak, and A. S. Reber. 1998. "Artificial Grammar Learning and the Mere Exposure Effect: Emotional Preference Tasks and the Implicit Learning Process." In *Handbook of Implicit Learning*, edited by M. A. Stadler and P. A. Frensch, 201–22. Thousand Oaks: SAGE PublicationsSage CA: Los Angeles, CA.

- Margolis, J. 1998. "Farewell to Danto and Goodman." *The British Journal of Aesthetics* 38 (4): 353–74.
- Margolis, J. 2000. "A Closer Look at Danto's Account of Art and Perception." *The British Journal of Aesthetics* 40 (3): 326–39.
- Markant, Julie, and Lisa S. Scott. 2018. "Attention and Perceptual Learning Interact in the Development of the Other-Race Effect." *Current Directions in Psychological Science* 27 (3): 163–69.
- Markie, Peter J. 2005. "The Mystery of Direct Perceptual Justification." *Philosophical Studies* 126 (3): 347–73.
- . 2006. "Epistemically Appropriate Perceptual Belief." *Nous* 40 (1): 118–42.
- Martindale, Colin, and Kathleen Moore. 1988. "Priming, Prototypicality, and Preference." *Journal of Experimental Psychology: Human Perception and Performance* 14 (4): 661–70.
- Mason, Elinor. 2015. "Moral Ignorance and Blameworthiness." *Philosophical Studies* 172 (11): 3037–57.
- Masrour, Farid. 2011. "Is Perceptual Phenomenology Thin?" *Philosophy and Phenomenological Research* 83 (2): 366–97.
- Matthen, Mohan. 2005. "Visual Concepts." *Philosophical Topics* 33 (1): 207–33.
- Matthews, Nestor, Zili Liu, Bard J. Geesaman, and Ning Qian. 1999. "Perceptual Learning on Orientation and Direction Discrimination." *Vision Research* 39 (22): 3692–3701.
- McCauley, Robert N., and Joseph Henrich. 2006. "Susceptibility to the Müller-Lyer Illusion, Theory-Neutral Observation, and the Diachronic Penetrability of the Visual Input System." *Philosophical Psychology* 19 (1): 79–101.
- McDowell, John. 1982. "Criteria, Defeasibility, and Knowledge." *Proceedings of the British Academy* 68: 455–79.
- . 1983. "Aesthetic Value Objectivity, and the Fabric of the World." In *Pleasure, Preference and Value*, edited by Eva Schaper, 1–16. Cambridge: Cambridge University Press.
- McElree, B, and M Carrasco. 1999. "The Temporal Dynamics of Visual Search: Evidence for Parallel Processing in Feature and Conjunction Searches." *Journal of Experimental Psychology*: 25 (6): 1517–39.
- McGovern, D. P., N. W. Roach, and B. S. Webb. 2012. "Perceptual Learning Reconfigures the Effects of Visual Adaptation." *Journal of Neuroscience* 32 (39): 13621–29.

- McGrath, Matthew. 2013. "Phenomenal Conservatism and Cognitive Penetration: The 'Bad Basis' Counterexamples." In *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, edited by Chris Tucker, 225–47. Oxford: Oxford University Press.
- . 2017. "Knowing What Things Look Like." *Philosophical Review* 126 (1): 1–41.
- McGugin, Rankin W., James W. Tanaka, Sophie Lebrecht, Michael J. Tarr, and Isabel Gauthier. 2011. "Race-Specific Perceptual Discrimination Improvement Following Short Individuation Training With Faces." *Cognitive Science* 35 (2): 330–47.
- McKee, Suzanne P., and Gerald Westhe. 1978. "Improvement in Vernier Acuity with Practice." *Perception & Psychophysics* 24 (3): 258–62.
- Mekawi, Yara, and Konrad Bresin. 2015. "Is the Evidence from Racial Bias Shooting Task Studies a Smoking Gun? Results from a Meta-Analysis." *Journal of Experimental Social Psychology* 61 (November): 120–30.
- Mervis, C B, and E Rosch. 1981. "Categorization of Natural Objects." *Annual Review of Psychology* 32 (1): 89–115.
- Michel, Caroline, Roberto Caldara, and Bruno Rossion. 2006. "Same-Race Faces Are Perceived More Holistically than Other-Race Faces." *Visual Cognition* 14 (1): 55–73.
- Mole, Christopher. 2015. "Attention and Cognitive Penetration." In *The Cognitive Penetrability of Perception: New Philosophical Perspectives*, edited by John Zeimbekis and Athanassios Raftopoulos, 218–37. Oxford University Press.
- . 2016. "Real Objective Beauty." *The British Journal of Aesthetics* 56 (4): 367–81.
- Montero, Barbara Gail. 2016. *Thought in Action: Expertise and the Conscious Mind*. Oxford: Oxford University Press.
- Morton, Adam. 2010. "Epistemic Emotions." In *The Oxford Handbook of Philosophy of Emotion*, edited by Peter Goldie, 385–99. Oxford: Oxford University Press.
- Mouchetant-Rostaing, Yolande, and Marie-Hélène Giard. 2003. "Electrophysiological Correlates of Age and Gender Perception on Human Faces." *Journal of Cognitive Neuroscience* 15 (6): 900–910.
- Mouchetant-Rostaing, Yolande, Marie-Hélène Giard, Shlomo Bentin, Pierre-Emmanuel Aguera, and Jacques Pernier. 2000. "Neurophysiological Correlates of Face Gender Processing in Humans." *European Journal of Neuroscience* 12 (1): 303–10.
- Mouchetant-Rostaing, Yolande, Marie-Hélène Giard, Claude Delpuech, Jean-Francois Echallier, and Jacques Pernier. 2000. "Early Signs of Visual Categorization for Biological and Non-Biological Stimuli in Humans." *Neuroreport*, no. 11: 2521–25.

- Murphy, Gregory L., and Hiram H. Brownell. 1985. "Category Differentiation in Object Recognition: Typicality Constraints on the Basic Category Advantage." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 11 (1): 70–84.
- Murphy, Gregory L., and EJ Wisniewski. 1989. "Categorizing Objects in Isolation and in Scenes: What a Superordinate Is Good For." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15 (4): 572.
- Mussweiler, Thomas. 2001. "'Seek and Ye Shall Find': Antecedents of Assimilation and Contrast in Social Comparison." *European Journal of Social Psychology* 31 (5): 499–509.
- Nakamura, K., A Mikami, and K Kubota. 1992. "Activity of Single Neurons in the Monkey Amygdala during Performance of a Visual Discrimination Task." *Journal of Neurophysiology* 67 (6): 1447–63.
- Nanay, Bence. 2015. *Aesthetics as Philosophy of Perception*. OUP Oxford.
- Neisser, Ulric, Robert Novick, and Robert Lazar. 1963. "Searching for Ten Targets Simultaneously." *Perceptual and Motor Skills* 17 (3): 955–61.
- Newell, Ben R., and James E.H. Bright. 2001. "The Relationship between the Structural Mere Exposure Effect and the Implicit Learning Process." *The Quarterly Journal of Experimental Psychology Section A* 54 (4): 1087–1104.
- Newell, Fiona N., and Heinrich H. Bülthoff. 2002. "Categorical Perception of Familiar Objects." *Cognition* 85 (2): 113–43.
- Newen, Albert, and Petra Vetter. 2017. "Why Cognitive Penetration of Our Perceptual Experience Is Still the Most Plausible Account." *Consciousness and Cognition* 47: 26–37.
- Nicki, Richard M., Poh Lin Lee, and Virginia Moss. 1981. "Ambiguity, Cubist Works of Art, and Preference." *Acta Psychologica* 49 (1): 27–41.
- Nodine, Calvin F., Harold L. Kundel, Claudia Mello-Thoms, Susan P. Weinstein, Susan G. Orel, Daniel C. Sullivan, and Emily F. Conant. 1999. "How Experience and Training Influence Mammography Expertise." *Academic Radiology* 6 (10): 575–85.
- O'Callaghan, Casey. 2011. "Against Hearing Meanings." *The Philosophical Quarterly* 61 (245): 783–807.
- O'Toole, Alice J, and Daniel J Kersten. 1992. "Learning to See Random-Dot Stereograms." *Perception* 21 (2): 227–43.
- Öhman, Arne, and Susan Mineka. 2001. "Fears, Phobias, and Preparedness: Toward an Evolved Module of Fear and Fear Learning." *Psychological Review* 108 (3): 483–522.
- Oliva, Aude, and Antonio Torralba. 2001. "Modeling the Shape of the Scene: A Holistic

- Representation of the Spatial Envelope.” *International Journal of Computer Vision* 42 (3).
- Osherson, Daniel N., and Edward E. Smith. 1981. “On the Adequacy of Prototype Theory as a Theory of Concepts.” *Cognition* 9 (1): 35–58.
- Palmeri, Thomas J., and Garrison W. Cottrell. 2010. “Modeling Perceptual Expertise.” In *Perceptual Expertise: Bridging Brain and Behavior*. Oxford University Press.
- Palumbo, Rocco, Bruno Laeng, and Luca Tommasi. 2013. “Gender-Specific Aftereffects Following Adaptation to Silhouettes of Human Bodies.” *Visual Cognition* 21 (1): 1–12.
- Pautz, A. 2009. “What Are the Contents of Experiences?” *The Philosophical Quarterly* 59 (236): 483–507.
- Peron, R M, and G L Allen. 1988. “Attempts to Train Novices for Beer Flavor Discrimination: A Matter of Taste.” *The Journal of General Psychology* 115 (4): 403–18.
- Pettit, Philip. 1983. “The Possibility of Aesthetic Realism.” In *Pleasure, Preference, and Value*, edited by Eva Schaper. Cambridge: Cambridge University Press.
- Pinker, Steven, and Alan Prince. 1999. “The Nature of Human Concepts: Evidence from an Unusual Source.” In *The Nature of Concepts : Evolution, Structure and Representation*, edited by Philip R. Van Loocke, 8–50. London: Routledge.
- Pisoni, David B., Richard N. Aslin, Alan J. Perey, and Beth L. Hennessy. 1982. “Some Effects of Laboratory Training on Identification and Discrimination of Voicing Contrasts in Stop Consonants.” *Journal of Experimental Psychology: Human Perception and Performance* 8 (2): 297–314.
- Plantinga, Alvin. 1993. *Warrant and Proper Function*. Oxford: Oxford University Press.
- Pollak, Seth D, Michael Messner, Doris J Kistler, and Jeffrey F Cohn. 2009. “Development of Perceptual Expertise in Emotion Recognition.” *Cognition* 110 (2): 242–47.
- Pollak, Seth D, and Pawan Sinha. 2002. “Effects of Early Experience on Children’s Recognition of Facial Displays of Emotion.” *Developmental Psychology* 38 (5): 784–91.
- Posner, Michael I. 1980. “Orienting of Attention.” *Quarterly Journal of Experimental Psychology* 32 (1): 3–25.
- Potter, Mary C., and Ellen I. Levy. 1969. “Recognition Memory for a Rapid Sequence of Pictures.” *Journal of Experimental Psychology* 81 (1): 10–15.
- Price, Richard. 2009. “Aspect-switching and Visual Phenomenal Character.” *The Philosophical Quarterly* 59 (236): 508–18.
- Prinz, Jesse. 2006. “Beyond Appearances: The Content of Sensation and Perception.” In

- Perceptual Experience*, edited by John P. Hawthorne, 434–60. Oxford University Press on Demand.
- . 2012. *The Conscious Brain*. Oxford University Press.
- . 2013. “Siegel’s Get Rich Quick Scheme.” *Philosophical Studies* 163 (3): 827–35.
- Pryor, James. 2000. “The Skeptic and the Dogmatist.” *Nous* 34 (4): 517–49.
- . 2005. “There Is Immediate Justification.” In *Contemporary Debates in Epistemology*, edited by M Steup and Ernest Sosa, 181–202. Malden, MA: Blackwell Publishing.
- Pylyshyn, Zenon W. 1984. *Computation and Cognition*. Cambridge, Massachusetts: MIT Press.
- . 1999. “Is Vision Continuous with Cognition?: The Case for Cognitive Impenetrability of Visual Perception.” *Behavioral and Brain Sciences* 22 (3): 341–65.
- . 2003. *Seeing and Visualizing : It’s Not What You Think*.
- Raftopoulos, Athanassios. 2001a. “Perceptual Learning Meets Philosophy: Cognitive Penetrability of Perception and Its Philosophical Implications.” In *Proceedings of the 23rd Annual Conference of the Cognitive Science Society*, edited by J. D. Moore and K. Stenning, 802–8. Mahwah, NJ: Lawrence Erlbaum.
- . 2001b. “Is Perception Informationally Encapsulated? The Issue of the Theory-Ladenness of Perception.” *Cognitive Science* 25 (3): 423–51.
- . 2009. *Cognition and Perception*. MIT Press.
- Rakison, David H, and Yevdokiya Yermolayeva. 2010. “Infant Categorization.” *Cognitive Science* 1: 894–905.
- Ramachandran, V S, and O Braddick. 1973. “Orientation-Specific Learning in Stereopsis.” *Perception* 2 (3): 371–76.
- Ramsey, Jennifer L., Judith H. Langlois, Rebecca A. Hoss, Adam J. Rubenstein, and Angela M. Griffin. 2004. “Origins of a Stereotype: Categorization of Facial Attractiveness by 6-Month-Old Infants.” *Developmental Science* 7 (2): 201–11.
- Ransom, Madeleine. 2019. “Frauds, Posers And Sheep: A Virtue Theoretic Solution To The Acquaintance Debate.” *Philosophy and Phenomenological Research* 98 (2): 417–34.
- . 2020. “Waltonian Perceptualism.” *Journal of Aesthetics and Art Criticism* 78 (1): 66–70.
- Ransom, Madeleine, Sina Fazelpour, and Christopher Mole. 2017. “Attention in the Predictive Mind.” *Consciousness and Cognition* 47.

- Raymond, Jennifer. 2013. "Most of Us Are Biased." *Nature* 495 (7439): 33–34.
- Reber, Rolf. 2002. "Reasons for the Preference for Symmetry." *Behavioral and Brain Sciences* 25 (3): 415–16.
- Reber, Rolf, and Norbert Schwarz. 2006. "Perceptual Fluency, Preference, and Evolution." *Polish Psychological Bulletin* 37 (1): 16–22.
- Reber, Rolf, Norbert Schwarz, and Piotr Winkielman. 2004. "Processing Fluency and Aesthetic Pleasure: Is Beauty in the Perceiver's Processing Experience?" *Personality and Social Psychology Review* 8 (4): 364–82.
- Reiland, Indrek. 2014. "On Experiencing High-Level Properties." *American Philosophical Quarterly* 51 (3): 177–87.
- Repp, Bruno H. 1997. "The Aesthetic Quality of a Quantitatively Average Music Performance: Two Preliminary Experiments." *Music Perception: An Interdisciplinary Journal* 14 (4): 419–44.
- Rhodes, Gillian, Susan Brake, Karyn Taylor, and Shirley Tan. 1989. "Expertise and Configural Coding in Face Recognition." *British Journal of Psychology* 80 (3): 313–31.
- Rhodes, Gillian, Keren Geddes, Linda Jeffery, Suzanne Dziurawiec, and Alison Clark. 2002. "Are Average and Symmetric Faces Attractive to Infants? Discrimination and Looking Preferences." *Perception* 31 (3): 315–21.
- Rhodes, Gillian, Linda Jeffery, Tamara L. Watson, Colin W.G. Clifford, and Ken Nakayama. 2003. "Fitting the Mind to the World: Face Adaptation and Attractiveness Aftereffects." *Psychological Science* 14 (6): 558–66.
- Rhodes, Gillian, Fiona Proffitt, Jonathon M. Grady, and Alex Sumich. 1998. "Facial Symmetry and the Perception of Beauty." *Psychonomic Bulletin & Review* 5 (4): 659–69.
- Rhodes, Gillian, Alex Sumich, and Graham Byatt. 1999. "Are Average Facial Configurations Attractive Only Because of Their Symmetry?" *Psychological Science* 10 (1): 52–58.
- Rhodes, Gillian, and Tanya Tremewan. 1996. "Averageness, Exaggeration, and Facial Attractiveness." *Psychological Science* 7 (2): 105–10.
- Rhodes, Marjorie, Sarah-Jane Leslie, and Christina M Tworek. 2012. "Cultural Transmission of Social Essentialism." *Proceedings of the National Academy of Sciences of the United States of America* 109 (34): 13526–31.
- Richler, Jennifer J., Yetta K. Wong, and Isabel Gauthier. 2011. "Perceptual Expertise as a Shift from Strategic Interference to Automatic Holistic Processing." *Current Directions in Psychological Science* 20 (2): 129–34.

- Rosch, Eleanor. 1973. "Natural Categories." *Cognitive Psychology* 4 (3): 328–50.
- Rosch, Eleanor, Carolyn B. Mervis, Wayne D. Gray, David M. Johnson, and Penny Boyes-Braem. 1976. "Basic Objects in Natural Categories." *Cognitive Psychology*.
- Rosch, Eleanor, and Carolyn B. Mervis. 1975. "Family Resemblances: Studies in the Internal Structure of Categories." *Cognitive Psychology* 7 (4): 573–605.
- Rossion, B., I. Gauthier, V. Goffaux, M.J. Tarr, and M. Crommelinck. 2002. "Expertise Training with Novel Objects Leads to Left-Lateralized Facelike Electrophysiological Responses." *Psychological Science* 13 (3): 250–57.
- Rotshtein, Pia, Richard N. A. Henson, Alessandro Treves, Jon Driver, and Raymond J. Dolan. 2005. "Morphing Marilyn into Maggie Dissociates Physical and Identity Face Representations in the Brain." *Nature Neuroscience* 8 (1): 107–13.
- Rousselet, Guillaume A., Michèle Fabre-Thorpe, and Simon J. Thorpe. 2002. "Parallel Processing in High-Level Categorization of Natural Images." *Nature Neuroscience* 5 (7): 629–30.
- Rousselet, Guillaume A., J. S. Husk, P. J. Bennett, and A. B. Sekuler. 2008. "Time Course and Robustness of ERP Object and Face Differences." *Journal of Vision* 8 (12): 1–18.
- Rousselet, Guillaume A., Marc J.-M. Macé, and Michèle Fabre-Thorpe. 2003. "Is It an Animal? Is It a Human Face? Fast Processing in Upright and Inverted Natural Scenes." *Journal of Vision* 3 (6): 5.
- Saito, Yuriko. 2001. "Everyday Aesthetics." *Philosophy and Literature* 25 (1): 87–95.
- . 2007. *Everyday Aesthetics*. Oxford: Oxford University Press.
- Samuels, Curtis A., George Butterworth, Tony Roberts, Lida Graupner, and Graham Hole. 1994. "Facial Aesthetics: Babies Prefer Attractiveness to Symmetry." *Perception* 23 (7): 823–31.
- Sangrigoli, S., C. Pallier, A.-M. Argenti, V.A.G. Ventureyra, and S. de Schonen. 2005. "Reversibility of the Other-Race Effect in Face Recognition During Childhood." *Psychological Science* 16 (6): 440–44.
- Saul, Jennifer. 2013. "Scepticism and Implicit Bias." *Disputatio* 5 (37): 243–63.
- Schweinberger, Stefan R., Romi Zäske, Christian Walther, Jessika Golle, Gyula Kovács, and Holger Wiese. 2010. "Young without Plastic Surgery: Perceptual Adaptation to the Age of Female and Male Faces." *Vision Research* 50 (23): 2570–76.
- Scott, Lisa S., James W. Tanaka, David L. Sheinberg, and Tim Curran. 2006. "A Reevaluation of the Electrophysiological Correlates of Expert Object Processing." *Journal of Cognitive Neuroscience* 18 (9): 1453–65.

- . 2008. “The Role of Category Learning in the Acquisition and Retention of Perceptual Expertise: A Behavioral and Neurophysiological Study.” *Brain Research* 1210 (May): 204–15.
- Seeck, M, C M. Michel, N Mainwaring, R Cosgrove, H Blume, J Ives, T Landis, and D L. Schomer. 1997. “Evidence for Rapid Face Recognition from Human Scalp and Intracranial Electrodes.” *NeuroReport* 8 (12): 2749–54.
- Segrest Purkiss, Sharon L., Pamela L. Perrewé, Trenea L. Gillespie, Bronston T. Mayes, and Gerald R. Ferris. 2006. “Implicit Sources of Bias in Employment Interview Judgments and Decisions.” *Organizational Behavior and Human Decision Processes* 101 (2): 152–67.
- Seitz, Aaron R., and Takeo Watanabe. 2009. “The Phenomenon of Task-Irrelevant Perceptual Learning.” *Vision Research* 49 (21): 2604–10.
- Shen, Li, Laurie R. Margolies, Joseph H. Rothstein, Eugene Fluder, Russell McBride, and Weiva Sieh. 2019. “Deep Learning to Improve Breast Cancer Detection on Screening Mammography.” *Nature: Scientific Reports* 9 (1): 12495.
- Sibley, Frank. 1959. “Aesthetic Concepts.” *The Philosophical Review* 68 (4): 421–50.
- . 2001a. “Aesthetic and Non-Aesthetic.” In *Approach to Aesthetics: Collected Papers on Philosophical Aesthetics*, edited by John Benson, Betty Redfern, and Jeremy Roxbee Cox, 33–51. Oxford: Clarendon Press.
- . 2001b. “Aesthetic Concepts.” In *Approach to Aesthetics: Collected Papers on Philosophical Aesthetics*, edited by John Benson, Betty Redfern, and Jeremy Roxbee Cox, 1–23. Oxford: Clarendon Press.
- . 2001c. *Approach to Aesthetics: Collected Papers on Philosophical Aesthetics*. Edited by John Benson, Betty Redfern, and Jeremy Roxbee Cox. Oxford: Clarendon Press.
- Siegel, Susanna. 2006a. “Which Properties Are Represented in Perception?” In *Perceptual Experience*, edited by Tamar Gendler and John Hawthorne, 481–503. Oxford: Oxford University Press.
- . 2006b. “Subject and Object in the Contents of Visual Experience.” *Philosophical Review* 115 (3): 355–88.
- . 2009. “The Visual Experience of Causation.” *The Philosophical Quarterly* 59 (236): 519–40.
- . 2010. *The Contents of Visual Experience*. New York: Oxford University Press.
- . 2012. “Cognitive Penetrability and Perceptual Justification*.” *Noûs* 46 (2): 201–22.
- . 2013. “The Epistemic Impact of the Etiology of Experience.” *Philosophical Studies* 162

- (3): 697–722.
- . 2017. *The Rationality of Perception*. Oxford: Oxford University Press.
- Siewert, C.P. 1998. *The Significance of Consciousness*. Princeton: Princeton University Press.
- Silins, Nicholas. 2013. “The Significance of High-Level Content.” *Philosophical Studies* 162 (1): 13–33.
- Smith, Brent W., and Michael B. Slack. 2015. “The Effect of Cognitive Debiasing Training among Family Medicine Residents.” *Diagnosis* 2 (2): 117–21.
- Smith, Edward E., and D.L. Medin. 1981. *Categories and Concepts*. Cambridge, Massachusetts: Harvard University Press.
- Smith, Steven M., Veronica Stinson, and Matthew A. Prosser. 2004. “Do They All Look Alike? An Exploration of Decision-Making Strategies in Cross-Race Facial Identifications.” *Canadian Journal of Behavioural Science / Revue Canadienne Des Sciences Du Comportement* 36 (2): 146–54.
- Sollberger, B., and Rolf Reber. 2004. “Artificial Grammar Learning as a Potential Mechanism for the Acquisition of Music Preference.” *Unpublished Manuscript*.
- Solomon, Gregg Eric Arn. 1990. “Psychology of Novice and Expert Wine Talk.” *The American Journal of Psychology* 103 (4): 495–517.
- Sowden, Paul T., Ian R. L. Davies, and Penny Roling. 2000. “Perceptual Learning of the Detection of Features in X-Ray Images: A Functional Role for Improvements in Adults’ Visual Sensitivity?” *Journal of Experimental Psychology: Human Perception and Performance* 26 (1): 379–90.
- Sowden, Paul T., Ian Davies, David Rose, and Martin Kaye. 1996. “Perceptual Learning of Stereoacuity.” *Perception* 25 (9): 1043–52.
- Spencer, Katherine B., Amanda K. Charbonneau, and Jack Glaser. 2016. “Implicit Bias and Policing.” *Social and Personality Psychology Compass* 10 (1): 50–63.
- Sperber, Dan. 2001. “In Defense of Massive Modularity.” In *Language, Brain, and Cognitive Development: Essays in Honor of Jacques Mehler*, edited by Emmanuel Dupoux, 47–58. Cambridge, Massachusetts: MIT Press.
- Srull, Thomas K., and Robert S. Wyer. 1979. “The Role of Category Accessibility in the Interpretation of Information about Persons: Some Determinants and Implications.” *Journal of Personality and Social Psychology* 37 (10): 1660–72.
- Stapel, Diederik A., and Jerry M. Suls, eds. 2007. *Assimilation and Contrast in Social Psychology*. Psychology Press.

- Stokes, Dustin. 2013. "Cognitive Penetrability of Perception." *Philosophy Compass* 8 (7): 646–63.
- . 2014. "Cognitive Penetration and the Perception of Art." *Dialectica* 68 (1): 1–34.
- . 2015. "Towards a Consequentialist Understanding of Cognitive Penetration." In *The Cognitive Penetrability of Perception : New Philosophical Perspectives*, edited by John Zeimbekis and Athanassios Raftopoulos, 75–102. Oxford: Oxford University Press.
- . 2018a. "Rich Perceptual Content and Aesthetic Properties." In *Evaluative Perception*, edited by A Bergqvist and Robert Cowan, 19–41. Oxford University Press.
- . 2018b. "Attention and the Cognitive Penetrability of Perception." *Australasian Journal of Philosophy* 96 (2): 303–18.
- . 2020. "On Perceptual Expertise." *Mind & Language*, mila.12270.
- Storrs, Katherine R. 2015. "Are High-Level Aftereffects Perceptual?" *Frontiers in Psychology* 6 (February): 157.
- Strahovnik, Vojko. 2004. "The Riddle of Aesthetic Principles." *Acta Analytica* 19 (33): 189–208.
- Subramaniam, Suresh, Irving Biederman, and Stephen Madigan. 2000. "Accurate Identification but No Priming and Chance Recognition Memory for Pictures in RSVP Sequences." *Visual Cognition* 7 (4): 511–35.
- Sugita, Yoichi. 2009. "Innate Face Processing." *Current Opinion in Neurobiology* 19 (1): 39–44.
- Susilo, T., E. McKone, and M. Edwards. 2010. "Solving the Upside-down Puzzle: Why Do Upright and Inverted Face Aftereffects Look Alike?" *Journal of Vision* 10 (13): 1–1.
- Suzuki, Satoru, and Patrick Cavanagh. 1998. "A Shape-Contrast Effect for Briefly Presented Stimuli." *Journal of Experimental Psychology: Human Perception and Performance* 24 (5): 1315–41.
- Tanaka, James W., and Tim Curran. 2001. "A Neural Basis for Expert Object Recognition." *Psychological Science* 12 (1): 43–47.
- Tanaka, James W., Tim Curran, Albert L. Porterfield, and Daniel Collins. 2006. "Activation of Preexisting and Acquired Face Representations: The N250 Event-Related Potential as an Index of Face Familiarity." *Journal of Cognitive Neuroscience* 18 (9): 1488–97.
- Tanaka, James W., Tim Curran, and David L. Sheinberg. 2005. "The Training and Transfer of Real-World Perceptual Expertise." *Psychological Science* 16 (2): 145–51.
- Tanaka, James W., and Martha J. Farah. 1993. "Parts and Wholes in Face Recognition." *The Quarterly Journal of Experimental Psychology Section A* 46 (2): 225–45.

- Tanaka, James W, Markus Kiefer, and Cindy M Bukach. 2004. "A Holistic Account of the Own-Race Effect in Face Recognition: Evidence from a Cross-Cultural Study." *Cognition* 93 (1): B1–9.
- Tanaka, James W, and Marjorie Taylor. 1991. "Object Categories and Expertise: Is the Basic Level in the Eye of the Beholder?" *Cognitive Psychology* 23 (3): 457–82.
- Tappolet, Christine. 2016. *Emotions, Values, and Agency*. Oxford: Oxford University Press.
- Thorpe, Simon J., Denis Fize, and Catherine Marlot. 1996. "Speed of Processing in the Human Visual System." *Nature* 381 (6582): 520–22.
- Thorpe, Simon J., Karl R. Gegenfurtner, Michèle Fabre-Thorpe, and Heinrich H. Bülthoff. 2001. "Detection of Animals in Natural Images Using Far Peripheral Vision." *European Journal of Neuroscience* 14 (5): 869–76.
- Todd, Cain. 2014. "Why We Don't Perceive Aesthetic Properties." In *Mind, Value and Metaphysics: Philosophical Papers Dedicated to Kevin Mulligan, Volume 2*, edited by Anne Reboul, 105–16. Springer.
- Tormey, Alan. 1973. "Critical Judgments." *Theoria* 39 (1–3): 35–49.
- Treisman, Anne M. 1993. "The Perception of Features and Objects." In *Attention, Selection, Awareness, and Control*, edited by Alan Baddeley and Lawrence Weiskrantz, 6–34. Oxford: Clarendon Press.
- Treisman, Anne M., and Garry Gelade. 1980. "A Feature-Integration Theory of Attention." *Cognitive Psychology*.
- Tucker, Chris. 2010. "Why Open-Minded People Should Endorse Dogmatism." *Philosophical Perspectives* 24 (1): 529–45.
- Tye, Michael. 2000. *Consciousness, Color, and Content*. MIT Press.
- Tye, Michael. 1995. *Ten Problems of Consciousness*. MIT Press.
- . 2002. "Representationalism and the Transparency of Experience." *Noûs* 36 (1): 137–51.
- Vaassen, Bram M.K. 2016. "Basic Beliefs and the Perceptual Learning Problem: A Substantial Challenge for Moderate Foundationalism." *Episteme* 13 (1): 133–49.
- Vance, Jona. 2014. "Emotion and the New Epistemic Challenge from Cognitive Penetrability." *Philosophical Studies* 169 (2): 257–83.
- VanRullen, Rufin. 2006. "On Second Glance: Still No High-Level Pop-out Effect for Faces." *Vision Research* 46 (18): 3017–27.

- VanRullen, Rufin, Lavanya Reddy, and Li Fei-Fei. 2005. "Binding Is a Local Problem for Natural Objects and Scenes." *Vision Research* 45 (25–26): 3133–44.
- VanRullen, Rufin, Lavanya Reddy, and Christof Koch. 2004. "Visual Search and Dual Tasks Reveal Two Distinct Attentional Resources." *Journal of Cognitive Neuroscience* 16 (1): 4–14.
- Vanrullen, Rufin, and Simon J. Thorpe. 2001. "The Time Course of Visual Processing : From Early Perception to Decision-Making .(Statistical Data Included) The Time Course of Visual Processing : From Early Perception to Decision-Making .(Statistical Data Included)." *Journal of Cognitive Neuroscience*, 1–7.
- VanRullen, Rufin, and Simon J. Thorpe. 2001. "The Time Course of Visual Processing: From Early Perception to Decision-Making." *Journal of Cognitive Neuroscience* 13 (4): 454–61.
- Vidyasagar, T R, and G W Stuart. 1993. "Perceptual Learning in Seeing Form from Motion." *Proceedings of the Royal Society of London. Series B: Biological Sciences* 254 (1341): 241–44.
- Vogt, Stine, and Svein Magnussen. 2007. "Expertise in Pictorial Perception: Eye-Movement Patterns and Visual Memory in Artists and Laymen." *Perception* 36 (1): 91–100.
- Walton, Kendall L. 1970. "Categories of Art." *The Philosophical Review* 79 (3): 334–67.
- . 2020. "Response to Commentators." *Journal of Aesthetics and Art Criticism* 1 (78): TBD.
- Watanabe, T., J. E. Nanez, and Y. Sasaki. 2010. "Perceptual Learning without Perception." *Journal of Vision* 1 (3): 467–467.
- Watanabe, T, J E Nández, and Y Sasaki. 2001. "Perceptual Learning without Perception." *Nature* 413 (6858): 844–48.
- Weaver, Matthew. 2018. "12 Camels Disqualified from Saudi Beauty Contest in 'Botox' Row." *The Guardian*, January 24, 2018.
- Webster, Michael A. 2011. "Adaptation and Visual Coding." *Journal of Vision* 11 (5): 3–3.
- . 2012. "Evolving Concepts of Sensory Adaptation." *F1000 Biology Reports* 4 (November): 1–7.
- . 2015. "Visual Adaptation." *Annual Review of Vision Science* 1 (1): 547–67.
- Webster, Michael A., Daniel Kaping, Yoko Mizokami, and Paul Duhamel. 2004. "Adaptation to Natural Facial Categories." *Nature* 428 (6982): 557–61.
- Webster, Michael A., and Donald I.A. Macleod. 2011. "Visual Adaptation and Face Perception."

- Philosophical Transactions of the Royal Society B: Biological Sciences* 366 (1571): 1702–25.
- Webster, Michael A., and Otto H. Maclin. 1999. “Figural Aftereffects in the Perception of Faces.” *Psychonomic Bulletin & Review* 6 (4): 647–53.
- Weisberg, R.W. 1999. “Creativity and Knowledge.” In *Handbook of Creativity*, edited by R.J. Sternberg, 226–50. Cambridge: Cambridge University Press.
- Werner, P.J. 2016. “Moral Perception and the Contents of Experience.” *Journal of Moral Philosophy* 13 (3): 294–317.
- Westheimer, Gerald, and Gert Hauske. 1975. “Temporal and Spatial Interference with Vernier Acuity.” *Vision Research* 15 (10): 1137–41.
- Westheimer, Gerald, and Suzanne P. McKee. 1978. “Stereoscopic Acuity for Moving Retinal Images*.” *Journal of the Optical Society of America* 68 (4): 450–55.
- Whitfield, T. W. A., and P. E. Slatter. 1979. “The Effects of Categorization and Prototypicality on Aesthetic Choice in a Furniture Selection Task.” *British Journal of Psychology* 70 (1): 65–75.
- Wiener, Richard L., Audrey T. Feldman Wiener, and Thomas Grisso. 1989. “Empathy and Biased Assimilation of Testimonies in Cases of Alleged Rape.” *Law and Human Behavior* 13 (4): 343–55.
- Wiesel, T N, and D H Hubel. 1963. “Single-Cell Responses in Striate Cortex of Kittens Deprived of Vision in One Eye.” *Journal of Neurophysiology* 26 (6): 1003–17.
- Wilson, C. L., T. L. Babb, E. Halgren, and P. H. Crandall. 1983. “Visual Receptive Fields And Response Properties Of Neurons In Human Temporal Lobe And Visual Pathways.” *Brain* 106 (2): 473–502.
- Winkielman, Piotr, Norbert Schwarz, Tedra A. Fazendeiro, and Rolf Reber. 2003. “The Hedonic Marking of Processing Fluency: Implications for Evaluative Judgment.” In *The Psychology of Evaluation: Affective Processes in Cognition and Emotion*, edited by Jochen Musch and Karl C. Klauer. London: Lawrence Erlbaum.
- Wolfe, Jeremy M., Kyle R. Cave, and Susan L. Franzel. 1989. “Guided Search: An Alternative to the Feature Integration Model for Visual Search.” *Journal of Experimental Psychology: Human Perception and Performance* 15 (3): 419–33.
- Wollheim, Richard. 1970. *Art and Its Objects*. Penguin Books.
- Wong, Alan C.-N., Thomas J. Palmeri, and Isabel Gauthier. 2009. “Conditions for Facelike Expertise With Objects: Becoming a Ziggerin Expert—But Which Type?” *Psychological Science* 20 (9): 1108–17.

- Wong, Alan C.-N., Thomas J. Palmeri, Baxter P. Rogers, John C. Gore, and Isabel Gauthier. 2009. "Beyond Shape: How You Learn about Objects Affects How They Are Represented in Visual Cortex." Edited by Jan Lauwereyns. *PLoS ONE* 4 (12): e8405.
- Wong, Yetta K., J. R. Folstein, and I. Gauthier. 2011. "Task-Irrelevant Perceptual Expertise." *Journal of Vision* 11 (14): 3–3.
- Wood, Greg, Karen M. Knapp, Benjamin Rock, Chris Cousens, Carl Roobottom, and Mark R. Wilson. 2013. "Visual Expertise in Detecting and Diagnosing Skeletal Fractures." *Skeletal Radiology* 42 (2): 165–72.
- Wood, Noelle L., and Nelson Cowan. 1995. "The Cocktail Party Phenomenon Revisited: Attention and Memory in the Classic Selective Listening Procedure of Cherry (1953)." *Journal of Experimental Psychology: General* 124 (3): 243–62.
- Wright, Crispin. 2002. "(Anti-)Sceptics Simple and Subtle: G.E. Moore and John McDowell." *Philosophy and Phenomenological Research* 65 (2): 330–48.
- . 2004. "Warrant for Nothing (and Foundations for Free)?" *Aristotelian Society Supplementary Volume* 78 (1): 167–212.
- Wu, Rachel, Alison Gopnik, Daniel C. Richardson, and Natasha Z. Kirkham. 2011. "Infants Learn about Objects from Statistics and People." *Developmental Psychology* 47 (5): 1220–29.
- Wu, Wayne. 2017. "Shaking Up the Mind's Ground Floor: The Cognitive Penetration of Visual Attention." *Journal of Philosophy* 114 (1): 5–32.
- Young, Andrew W, Deborah Hellawell, and Dennis C Hay. 2013. "Configurational Information in Face Perception." *Perception* 42 (11): 1166–78.
- Young, James O. 1997. "Relativism and the Evaluation of Art." *Journal of Aesthetic Education* 31 (1): 9–22.
- Young, Steven G., and Kurt Hugenberg. 2012. "Individuation Motivation and Face Experience Can Operate Jointly to Produce the Own-Race Bias." *Social Psychological and Personality Science* 3 (1): 80–87.
- Zangwill, Nick. 1998. "Aesthetic/Sensory Dependence." *British Journal of Aesthetics* 38 (1): 66–81.
- . 2001. *The Metaphysics of Beauty*. Ithaca: Cornell University Press.
- Zeimbekis, John, and Athanassios Raftopoulos. 2015a. "The Cognitive Penetrability of Perception: An Overview." In *The Cognitive Penetrability of Perception: New Philosophical Perspectives*, edited by John Zeimbekis and Athanassios Raftopoulos, 2–58. Oxford: Oxford University Press.

———. , eds. 2015b. *The Cognitive Penetrability of Perception: New Philosophical Perspectives*. Oxford University Press USA - OSO.