

Aesthetic Perception and the Puzzle of Training **Madeleine Ransom**

Abstract: While the view that we perceive aesthetic properties may seem intuitive, it has received little in the way of explicit defence. It also gives rise to a puzzle. The first strand of this puzzle is that we often cannot perceive aesthetic properties of artworks without training, yet much aesthetic training involves the acquisition of knowledge, such as when an artwork was made, and by whom. How, if at all, can this knowledge affect our perception of an artwork's aesthetic properties? The second strand of the puzzle arises when we widen the scope of aesthetic experience. The very same aesthetic properties that seem to require training for their perception in artworks do not appear require training to perceive in objects of everyday aesthetic appreciation and natural phenomena. In this paper I argue that a prominent extant attempt to explain how training is compatible with aesthetic perception – cognitive permeation – is an inadequate solution. I also develop a positive view of aesthetic perception that provides a unified solution to both strands of the puzzle.

Keywords: *aesthetic perception; cognitive penetration; perceptual learning; perceptual expertise; categories of art*

1. Introduction

When we watch a dance performance, stand before a painting or listen to music, we do not seem to reason our way to the judgment that the dancers are graceful, the painting is beautiful, or the music is soulful. Instead, these aesthetic properties appear to be an inextricable part of the fabric of our experience.¹ The view that we perceptually experience – or experience in some broader sense – aesthetic properties, is the predominant view in aesthetics (Hopkins, 2005; Irvin, 2008; Iseminger, 2004; Lamarque, 2010; Leopold, 1949; Levinson, 2005; Lopes, 2014a, 2016; Matey, 2016; Sibley, 2001b; Stokes, 2014; Tormey, 1973; Tucker, 1997; Wollheim, 1970; Zangwill, 1998). Yet it has received surprisingly little defence, and has recently been challenged on the grounds that the view is in tension with the training required to develop aesthetic expertise (Dorsch, 2013). In what follows, I address this challenge and defend a strong version of the view: we represent at least some aesthetic properties in perceptual experience (hereafter I refer to this view as 'aesthetic perception').²

Whether or not aesthetic perception is correct matters for our understanding of the epistemology of aesthetic beliefs. For example, if we perceive aesthetic properties, then this suggests our justification for aesthetic beliefs is epistemically immediate: these beliefs do not depend on other beliefs for their justification (Alston, 1983; though see Jenkin, 2020). The truth of aesthetic perception also matters to our understanding of the nature of perceptual experience. It is uncontroversial that we represent low-level properties such as shape, size, motion, and colour in perceptual experience. However, it is controversial whether our perceptual experience represents any high-level properties. High-level properties are usually defined stipulatively as any properties that go beyond the basic 'building blocks' of perceptual experience such as shape and colour. While the debate has largely focused on artifactual or natural kind properties, such as being a tea kettle or a tiger (Brogaard, 2013;

¹ Exactly what makes a property aesthetic is a vexed question (De Clercq, 2002), and beyond the scope of this paper. In what follows I simply rely on intuitions that the properties I discuss are aesthetic.

² A stronger thesis would be that we represent all aesthetic properties in perceptual experience. I do not defend this thesis here because it would require a principled way of identifying all those properties that count as aesthetic. The more liberal the account of aesthetic properties, the harder such a thesis would be to defend.

Siegel, 2010), at least some aesthetic properties are also plausibly viewed as high-level properties (Logue, 2013, 2018; Matey, 2016; Stokes, 2018).

One reason to believe this is that aesthetic properties are not generally taken to be wholly reducible to combinations of simple shapes, colours, and the like. Just as the high-level property ‘dog’ tracks a category that includes dogs of various shapes and sizes, so too do aesthetic properties such as ‘being garish’ or ‘elegant’ track categories that include artifacts and natural phenomena of various shapes and sizes. While individual instances of individual dogs or objects of aesthetic interest may be wholly describable in terms of low-level properties, the category is not a mere congerie of such properties (Ransom, 2020a).

A second, related, reason for taking aesthetic properties to be high-level properties is that at least some aesthetic properties are often held to possess an evaluative component (De Clercq, 2008). To ascribe the property of elegance, say, to an object is not just to describe it but also to endorse it (see for example Sibley, 2001c). This is another way of denying that aesthetic properties are wholly reducible to low-level properties. The upshot for the general debate over the contents of perceptual experience is as follows. If we perceive aesthetic properties then, insofar as aesthetic properties are high-level properties, this vindicates the view that high-level properties are sometimes part of the contents of perceptual experience.

One challenge to aesthetic perception is that aesthetic expertise seems to require a great deal of training. We are not born wine connoisseurs or experts on the works of female Renaissance painters. Instead, we acquire these abilities over time, through training that often involves formal instruction. Training does not sit well with the perceptual view for two reasons. First, if we suppose that aesthetic training essentially involves learning about the artist, historical milieu, production techniques used, and so on – what I shall call *art-historical knowledge* – then, absent an account of how such knowledge can alter perceptual experience, this favours the view that we instead draw upon this knowledge in a purely cognitive, inferential process in order to make aesthetic judgments.³ After all, acquiring art-historical knowledge plausibly allows people to make a wider range of inferences about the artwork (Cavedon-Taylor, 2017). While it may not seem to us as if we reason our way to aesthetic property attributions, it may be that we are engaging in implicit or internalized reasoning that nevertheless draws upon this knowledge (Dorsch, 2013).

Second, training does not sit well with aesthetic perception as it invites a disanalogy with our ability to perceive low-level properties such as size, colour, and shape. Consider the conception of training outlined above in terms of art-historical knowledge. The correlate with respect to low-level properties would be learning to perceive them through the acquisition of background knowledge, such as that squares are shapes that possess four sides of equal length connected at right angles. We do not seem to require such background knowledge to perceive shapes, though we may require it in order to form beliefs about our perceptual experience. So, if training is involved in the case of aesthetic properties – or, indeed, in the case of any high-level property – then this counts against a perceptual view absent a plausible explanation of how training alters perceptual experience.⁴ Explaining how aesthetic perception is compatible with training is what I will call the *challenge from aesthetic training*. It is also the first strand of a broader puzzle over the role of training in aesthetics.

³ Here I understand aesthetic judgments to be beliefs about what aesthetic properties an object possesses, and distinct from perceptual representations of aesthetic properties. The perceptual experience of aesthetic properties is one route by which we may arrive at aesthetic judgments, but we may also arrive via inference or testimony (see for example Ransom, 2019).

⁴ Dorsch (2013) takes the role of training to be a problem for all accounts that posit that we perceive learned high-level properties, beyond aesthetics, and so while here I focus on aesthetic properties, much of what follows is also applicable to other high-level properties.

The second strand of the puzzle stems from widening the scope of aesthetic investigation. If we restrict ourselves to the fine arts or highly specialized aesthetic practices, then the necessity of some sort of training for making accurate aesthetic judgments in a variety of cases seems plausible.⁵ In contrast, if we focus on everyday aesthetics or environmental aesthetics, then it is instead training that many will find implausible. Aesthetic experience is a common part of our lives, despite the apparent lack of training.

Everyday aesthetics emphasizes the pervasive nature of the aesthetic. Our experience of aesthetic properties is part of the fabric of our lives, from choosing what clothing to put on in the morning, to selecting a tea cup (Irvin, 2008; Saito, 2001). While these aesthetic experiences are not as intense and focused as our experiences with symphonies and artworks in museums, they are nevertheless important sources of value. There is of course art-historical knowledge to be had about fashion and ceramics. However, there is a sense in which our everyday experience appears to put us in touch with aesthetic properties without requiring such knowledge. We can have a favourite mug, and admire its heft, delicate pattern, or the way it fits perfectly in our hands without knowing anything about ceramic production techniques or historical trends.

There are also arguably cases of everyday aesthetic experience where there is no such art-historical knowledge to be had. Some of the examples Irvin (2008) provides are smelling her cat's fur, scratching an itch, and running her tongue over her teeth to feel their smoothness. If such cases count as aesthetic, then it is unclear what sort of background knowledge could be involved in training.

Appreciation of nature is another case, albeit more controversial, as some hold that background knowledge in the form of scientific knowledge of an object's category or function is required in this domain as well (Carlson, 1979; Eaton, 1998; Parsons & Carlson, 2004). If it is the case that we require such knowledge for aesthetic appreciation then it will be an integral part of aesthetic training, and the second strand of the puzzle dissolves. This would not be a dissolution in favour of aesthetic perception, however. Defenders of the view would now also be required to shoulder the burden of explaining how background knowledge of the natural world alters our perceptual experience in order to represent aesthetic properties. And here there is no guarantee that the account offered in the case of art-historical knowledge will also generalize to knowledge of the natural world.

However, others hold the intuition that no such background knowledge is required (Brady, 1998; Carroll, 1993; Walton, 1970; Zangwill, 2001). We seemingly effortlessly appreciate the beauty of sunsets, the cuteness of small animals, and the ugliness of naked mole rats without background knowledge. This intuition is bolstered in the case of human natural beauty, with empirical studies showing that babies just a few months old prefer to look longer at more beautiful faces (Ramsey et al., 2004; Rhodes et al., 2002; Samuels et al., 1994). Babies have not plausibly undergone training, and yet have preferences that reliably track faces judged beautiful in a culture. While such empirical studies are not decisive, and the realm of nature appreciation extends far beyond human beauty, so long as one grants that there is at least some untrained aesthetic appreciation in some domains, then the second strand of the puzzle stands.⁶ One might, for example, hold that while some appreciation of

⁵ Here I do not mean to claim that training is required to perceive all aesthetic properties in all artworks. There are plausibly lots of 'untrained' aesthetic experiences when we engage with artworks as well. The challenge from aesthetic training is rather to explain how training and perception are compatible in cases where the ascription of aesthetic properties seems to require expertise. Thanks to an anonymous reviewer for pressing me to clarify this point.

⁶ See also (Dutton, 1979). While he is skeptical of untrained aesthetic experiences in adults, he nevertheless grants that babies likely have some sort of aesthetic experience. This minimal concession is sufficient to motivate the second strand of the puzzle.

nature depends on knowledge of the kind of thing that it is, one can also aesthetically appreciate natural objects simply by considering how the thing is in itself, apart from function or category (Zangwill, 2001).

Explaining untrained aesthetic appreciation does not pose a challenge to aesthetic perception per se. Lack of training even counts in favour of aesthetic perception insofar as in these cases there is no longer a disanalogy between low-level properties and aesthetic properties in terms of the training required. However, the contrast between the need for training and the lack of training does raise a challenge, call it the *challenge from aesthetic asymmetry*: If we can perceive aesthetic properties in some cases – such as in some natural objects – without training, then why can't we also perceive them in others – such as some artworks – without training? The challenge is to explain this asymmetry.

This challenge would be easily resolvable if aesthetic properties could be separated neatly into two groups: the aesthetic properties that don't require training to perceive and those that do. Then we might hold that we attribute aesthetic properties in the former category to natural and everyday objects, and those in the latter category to artworks. For example, we might hold that only artworks or artefacts can be witty, dull, or cloying. However, this is implausible as a general solution given there are aesthetic properties that seem to straddle both categories. For example, we may perceive a face as beautiful from the get-go, but only come to see the beauty of a Leonora Carrington painting after training. We may perceive the gracefulness of a swan's neck without a second thought, but only come to appreciate the gracefulness of an Indian classical dancer's performance with familiarity and knowledge of the genre.⁷

This asymmetry is a bizarre feature of aesthetic life, especially when we compare the situation with respect to low-level properties, or even other high-level properties. If we can perceive the colour red in flowers and clothing, then we should have no trouble perceiving the same colour in Malevich's *Red Square*. If we can perceive a crow perched on a telephone wire across the street, then it should be no problem perceiving the crows that populate many of Andrea Kowch's paintings. How can the same aesthetic properties strike us in some cases but elude us in others, if indeed they are perceptual?

The *puzzle of training* is how aesthetic perception can be both trained and untrained. It consists of two related challenges to any theory of aesthetic perception. The challenge from aesthetic training requires an account of how training is compatible with aesthetic perception, and the challenge from aesthetic asymmetry requires an account of why we don't need training to perceive aesthetic properties in some objects, but do need it to perceive the same properties in other objects. In the following section, I argue that cognitive permeation provides an inadequate solution to both strands of the puzzle of training.

2. Aesthetic cognitive permeation

Cognitive permeation – also commonly known as cognitive penetration – is hypothesized to occur when the contents of perceptual experience are altered by one's cognitive states such as beliefs and desires, where this alteration meets at least two

⁷ An anonymous reviewer for this journal suggests one might answer the challenge by holding that natural beauty and artefactual beauty are different aesthetic properties. While I don't have the space here to fully consider this solution, the aesthetic perceptualist who adopts this would have to provide a principled argument for construing them as different properties, beyond simply pointing to their artefactual/natural status, or else risk the solution being ad hoc. Moreover, many everyday objects are also artefacts, such as mugs and vases. If no training is required to perceive some aesthetic properties in these objects, then the challenge from aesthetic asymmetry remains.

conditions.⁸ The causal condition holds that the permeating belief must be a cause of the resulting perceptual experience. On the internal condition, the permeating belief must provide direct input to perceptual processing, which then modulates perceptual experience. A true belief that it's raining outside might cause me to go to the window and so change my perceptual experience to one of perceiving rain, but since this is mediated by my external action it would not count as cognitive permeation. Conversely, if this belief caused me to hallucinate rain where there is none, then this would count as cognitive permeation. Often a third condition – positing that the content of the permeating belief must bear a semantically non-arbitrary relationship to the resulting content of the perceptual experience – is also included, though exactly how to elaborate it has been the subject of significant debate (Pylyshyn, 1999; Stokes, 2013). In what follows I will not assume this third condition to be required, as none of my arguments turn on it.⁹

Several philosophers have proposed cognitive permeation as an explanation for how training in the form of art-historical knowledge can allow people to come to perceive both (i) artworks as belonging to a category of art, and (ii) aesthetic properties (Hopkins, 2005; Lamarque, 2010; Margolis, 1998; Nanay, 2015). I will focus on each of these proposed explanations in turn.

2.1 *Perceiving categories of art*

The thesis that we perceive artworks as belonging to categories of art was first proposed by Kendall Walton (1970). Dustin Stokes (2014) argues that cognitive permeation fills an explanatory gap in Walton's account, which he terms the *expertise-to-perception effect*: how training can lead to perceiving artworks in categories. That is, it answers the challenge from aesthetic training. Walton himself does not say much about this beyond that exposure to many artworks of a given category is a part of this training (1970, p. 366).

There are two components to Stokes' (2014) proposal for how we come to represent categories of art in perceptual experience. First, he takes art-historical knowledge along the lines of “*this* is the most critically praised Monet or *this* is an impressionist exhibit” (p. 12) to guide categorization. This guidance is explicit at first as one is beginning training. It then becomes largely implicit, as one develops increased expertise. Such guidance is supposed to amount to cognitive permeation because knowledge of this sort will then internally affect perceptual experience. On the liberal view Stokes proposes, knowledge of the category of art leads to our representing the artwork as belonging to that category (a kind of high-level property) in perceptual experience. We come to perceive artworks as exemplars of Impressionist paintings, Dogon masks, or *śāstriya saṅgīt* music.¹⁰

⁸ Here I follow Becko Copenhaver (personal conversation) in using cognitive permeation to avoid gendering reason as masculine.

⁹ Readers may wonder why I don't discuss the consequentialist account of cognitive permeation proposed by Dustin Stokes (2014), given that I focus on his view of aesthetic cognitive permeation in what follows. On the consequentialist view, for a cognitive-perceptual relation to count as cognitive permeation it must have consequences for the epistemic role of perception, modularity, or the theory-ladenness of perception. However this characterization is problematic, as the alternative to cognitive permeation I propose in section 3 also has the same epistemic consequence as the one Stokes (2014, p.13) outlines: people can gain more knowledge about artworks because they better perceive their aesthetic properties. Judged purely on consequences, my proposal would thereby be an instance of cognitive permeation. But there are good independent reasons to deny this – as I discuss further in section 3, perceptual learning does not require background knowledge to occur, and so in such cases there can be no theoretically significant causal, internal influence of cognition on perception of the sort hypothesized on even consequentialist accounts of cognitive permeation. Thanks to an anonymous reviewer for pressing me on this issue.

¹⁰ Stokes also considers a conservative view, where knowledge of artistic category alters our perceptual experience of aesthetic properties without the category itself being represented in perceptual experience. I take the liberal view to be closer to Walton's original view.

The second component of Stokes' proposal is that through training we come to learn new patterns of attention to the relevant features of the artworks: critics receive explicit training about where to look and what properties to attend to, and these patterns of attention are eventually internalized and automated, so that they are no longer under agential control. Once this has occurred, it counts as an instance of cognitive permeation because art-historical knowledge has systematically altered our perceptual experience. We learn to attend to aspects of the work that are partially constitutive of perceiving objects in categories of art.¹¹

This second component concords with Walton's (1970) hypothesis that what it is 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, such as the standard colour palette, subject matter, and brushstroke style of impressionist paintings. An additional condition on perceiving an artwork in a category is that there not be too many contrastandard properties present in the object. A contrastandard property is one that the work possesses that tends to disqualify it from membership in that category, though it may not be sufficient to do so on its own. For example, a very dark colour palette may count against a work's qualifying as an impressionist painting, though the brushstroke style and typical subject matter may nevertheless allow it to be categorized as such. 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.

While Stokes himself does not suggest this, one way of connecting his two proposals is that our internalized attentional patterns track Walton's standard properties. The automated attentional selection and grouping of these standard perceptible properties may be what explains how they come to be unified into a perceptual gestalt that corresponds to a high-level property such as 'being an Impressionist painting' (Ransom, 2020a).

However, perceiving artworks in categories does not get us all the way to the goal of explaining how we perceive aesthetic properties. There is still a second explanatory gap between perceiving an artwork in a category and perceiving its aesthetic properties – what Stokes calls the *perception-to-aesthetic reaction effect*.

2.2 *Perceiving aesthetic properties*

Those who claim that art-historical knowledge affects what categories of art we perceive a work as belonging to also tend to endorse Walton's (1970) psychological thesis that what category we perceive an artwork in alters what aesthetic properties the work appears to have. Walton defends this thesis by way of examples meant to demonstrate that perceiving non-aesthetic properties as standard, contrastandard or variable relative to a category in turn affects what aesthetic properties the work appears to have. Given that sculptures are standardly static, a sculpture with a twitching element – a contrastandard property – would be shocking, where 'shocking' is meant to be an aesthetic property. 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

¹¹ It is not entirely clear that Stokes means attention to explain our perception of categories of art rather than the perception of aesthetic properties themselves. Here I have interpreted him as intending the former, because he is committed to the claim that through perceiving artworks in categories we come to perceive aesthetic properties, as I elaborate in the next section.

exposition-development-recapitulation structure. Variable properties will also vary widely in how they factor into our perception of aesthetic properties, as they count neither for nor against a work's categorization.

The psychological thesis does not, however, give Walton the resources to say when someone is mistaken in their aesthetic judgment. If one person perceptually categorizes an artwork as a sculpture, and another person categorizes it as a painting, then for all Walton has said both may have accurate aesthetic experiences, or both may be a species of perceptual illusion. In order to secure accuracy conditions on perceiving aesthetic properties, Walton defends 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. That is, veridical categories depend in part on art-historical facts. While Walton also holds such categories will depend on which category allows the work to come off as aesthetically better, this will only be in the absence of any contravening art-historical facts.¹²

Walton's metaphysics of aesthetic properties is then as follows. He endorses Frank Sibley's thesis that aesthetic properties depend on an object's perceptible non-aesthetic properties. To Sibley's thesis he adds only that some of these non-aesthetic properties are standard, contrastandard, and variable properties, where these are metaphysically dependent on what category of art a work belongs to. Sibley conceived of aesthetic properties as emergent properties, where the non-aesthetic properties of an object determine its aesthetic properties (Sibley, 2001a). Walton adds to this account that an object's aesthetic properties will be determined in part by the category(ies) of art it belongs to, which is determined by art-historical facts (1970, pp. 337–338).

The idea is that once we allow that we perceive the artwork in a category then we get aesthetic perception for free. This is because to perceive an artwork in a category just is to perceive certain non-aesthetic properties as standard, contrastandard, and variable, and aesthetic properties will supervene on these properties. The perception-to-aesthetic reaction effect can be wholly explained by the acquisition of art-historical knowledge of categories of art, which are the only beliefs said to permeate perceptual experience in the manner sketched out in section 2.1. No separate explanation of how we can come to perceive aesthetic properties in artworks via training is required.

2.3 Evaluating the cognitive permeation hypothesis

Stokes argues that the cognitive permeation of perception provides a more unified explanation than the hypothesis that our aesthetic judgments are based on cognitive inferences. He claims that it concords with our pre-theoretical intuitions that we perceive aesthetic properties, and it also concords with the cognitive permeation interpretation of a large body of empirical work that does not pertain to aesthetic properties: if it is the best explanation of these studies, then it is the best explanation in the aesthetic case (2014, p22). This claim that cognitive permeation is the best explanation of a wide body of empirical results has recently come under scrutiny (Firestone & Scholl, 2014, 2016; Machery, 2015).

¹² The psychological and normative theses do not appear to sit happily together (Laetz, 2010). On the one hand, attributing aesthetic properties is something that happens in perception. On the other, such (correct) attributions depend on art-historical knowledge. How can knowledge influence perception? In the face of this apparent tension Stokes proposes that the normative thesis relates to the psychological thesis via cognitive permeation (see also Lamarque, 2010).

However, independently of this, cognitive permeation does not provide a unified explanation of how aesthetic perception occurs.

First, a necessary (but perhaps not sufficient) condition for cognitive permeation is that one possess the relevant art-historical knowledge, or else there is no candidate belief that causes the change in perceptual experience. This precludes the possibility that completely self-taught individuals who have not acquired such art-historical knowledge can come to perceive artworks in categories. However, Walton is explicit that art-historical knowledge is not required. He holds that the critic's ability to perceive an artwork in the correct category "does not require consideration of historical facts, or consideration of facts at all" (1970, 366). Moreover, in Walton's original introduction of the thesis that we can come to perceive artworks in perceptually distinguishable categories, he states that this excludes categories that require invoking non-perceptual considerations in order to determine correct membership. Stokes may also grant that such cases exist. He claims that some people are 'natural appreciators' that do not receive any formal artistic training in how to direct their attention (2014, pp.26-29).¹³

There is also some empirical evidence for categorization of art absent art-historical knowledge. In machine learning, neural networks can be trained to either categorize artworks by genre (Lecoutre et al., 2017) or individual artistic style (Seeley et al., 2017), based only on the labels and perceptible features of the training data. While artificial neural networks should not be assumed to rely on the same features that humans do to make categorizations, such studies nevertheless provide proof of concept that some categories of art are potentially learnable via perceptual cues alone, provided a large enough training set of images. What must be shown is that there are sufficient perceptual cues available to humans in order to be able to make such categorizations. Here, experimental work in comparative cognition provides evidence that species such as rats, birds, monkeys, and fish can distinguish between different categories of art, such as classical and blues music (Watanabe, 2012). The similarity of the human perceptual apparatus to that of at least some of these species points to the availability of sufficient perceptual cues for humans (see also Bullot & Reber, 2013 section 3.1.1. for a discussion of some human empirical studies).

Given these considerations, we are left with a new explanatory gap: how it is that people can sometimes come to perceive artworks in categories without acquiring the requisite art-historical knowledge. As long as there are real cases of these autodidacts, cognitive permeation does not offer a complete explanation for the phenomenon of aesthetic perception. The solution must be combined with another explanation, and so does not offer a unified solution to the challenge from aesthetic training – what Stokes terms the expertise-to-perception effect.

Let us now turn to the second component of the view, how the cognitive permeation hypothesis explains the perception-to-aesthetic reaction effect; that is, how we come to perceive aesthetic properties themselves. Recall that the extent of the explanation is (i) the psychological thesis that what category of art we perceive an object in will affect what aesthetic properties we perceive artworks as having, and (ii) we will veridically perceive aesthetic properties when such categories are correct, as an object's aesthetic properties supervene on its non-aesthetic properties, including which properties are standard, contrastandard, or variable relative to the correct category. If this account is correct, then

¹³ While Stokes holds that natural appreciators will still acquire art-historical knowledge in their training, given that this knowledge will not overtly direct their attention it is not clear how cognitive permeation is supposed to occur in these instances.

perhaps we should take this supervenience to be explanatorily basic. No separate explanation of how we can come to perceive aesthetic properties in artworks via training is required.

This solution falters on two counts, however. First, it does not explain cases where aesthetic perception and subsequent aesthetic judgments go wrong in spite of perceiving artworks in the correct category. While Walton (1970) holds only the thesis that what aesthetic properties a work has will be *often* determined *in part* by which properties are (correctly) perceived as standard, contrastandard, or variable, he does not explain the additional factors that exert an influence (see for example 343, 354). The cognitive permeation hypothesis does not provide any additional insight. Moreover, defending the stronger thesis that which properties are standard, contrastandard, or variable completely determine which aesthetic properties a work has is implausible because it leaves no room for disagreement over what aesthetic properties a work has amongst two people who both perceive the artwork in the correct category.

Do such cases exist? An informal survey of aesthetic criticism suggests that they do. Those who watch television series *The 100* likely recognize that it is science fiction. The main premise is that a hundred juvenile delinquents are sent from a dying spaceship to determine whether a nuclear-holocaust ravaged Earth is again habitable. The show is pretty on-genre. Yet reviewers on [rottentomatoes.com](https://www.rottentomatoes.com) disagree on its aesthetic properties, with one audience member noting of season two, “Pretty interesting storyline but got boring fast” and another writing “It has cliffhangers on every episode and there is a lot of suspense/thrill on each episode!”¹⁴ Assuming that boring and suspenseful are aesthetic properties, and that they are in opposition to one another, then this is a qualifying case. Similar aesthetic disagreements can be found for other TV series and movies, as well as for music. People agree that Celine Dion makes pop music heavy on the power ballads, but they disagree on the aesthetic merits of her music (Wilson, 2007).

Second, claims of ontological emergence or supervenience merely state that a pattern holds, and fail to provide an explanation of why it holds. Yet, offering an explanation is often feasible in such cases (Horgan, 1993). We can and ought to do better in terms of offering a complete account of how we come to represent aesthetic properties in perceptual experience.

Here one way of explaining the phenomenon is that, just as we can form perceptual categories of art by acquiring background knowledge, 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 recognize being graceful as such – then one can perceptually represent, for example, a ballet dancer as graceful” (Stokes, 2014, p. 12). The explanation here is then that through the acquisition of background knowledge of what aesthetic properties are or the conditions for their application, we can come to represent aesthetic properties directly via cognitive permeation, independently of categories of art.

Such background knowledge is plausibly acquired either by inference, testimony, or perception.¹⁵ 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 one must deny aesthetic particularism, the thesis that there are no general

¹⁴ The first is a quote from Madisyn W (12/22/20) and the second from Seth B (1/18/21). Accessed October 26, 2021: https://www.rottentomatoes.com/tv/the_100/s02/reviews?type=user

¹⁵ There are more options here, such as a priori intuition. I leave it to others to explore whether these options are plausible.

rules for inferring aesthetic properties from non-aesthetic properties (Bender, 1995; Lopes, 2014b; Sibley, 1959; Strahovnik, 2004). Particularism is motivated by the actual absence in aesthetic training of such rules, and by the highly contextual nature of aesthetic properties. For a particular painting it might be true that a bright yellow spot makes it dynamic, but such a spot on another painting might have the effect of rendering it lifeless. Aesthetic properties are not plausibly like artifactual kind properties such as telephones. Objects that possess a common aesthetic property may vary widely in their other 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. Choosing this route requires finding and defending plausible general principles of inference, which is a seemingly Sisyphean task.

Appealing to inductive inference or testimony (Cavedon-Taylor, 2017) also does not help. I grant that our knowledge that certain aesthetic properties are instantiated is sometimes derived from these sources (Ransom, 2019). Based on our familiarity with a band's past repertoire, or the testimony of a trusted expert, we can form a reasoned belief that their next song is likely to be melancholic. But when we ask how we (or the expert) arrived at these past judgments, the answer again must be 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 that in turn allow us to form either beliefs that the aesthetic properties are instantiated or more general beliefs concerning their application conditions, and these beliefs then permeate our perceptual experience. This last step is not required, however, as the chain itself begins with perception of aesthetic properties.

To summarize, the cognitive permeation view is only a partial solution to the puzzle of training. First, it does not explain cases where no art-historical knowledge is required to perceive artworks in categories. Second, it doesn't explain how aesthetic perception occurs beyond a bare supervenience account. Third, it does not address the challenge from aesthetic asymmetry: it does not explain how we are able to come to perceive aesthetic properties in cases such as everyday and natural objects. To be fair, the account does not purport to explain cases where no background knowledge is required for aesthetic perception. However, this signifies that a separate account of aesthetic perception that does not invoke cognitive permeation must also be provided. Then the issue is whether this separate account may serve as an explanation of the cases that, at first glance, seem to require cognitive permeation.

In what follows I propose an alternate account of aesthetic perception that provides a unified solution to the puzzle of training. 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. Therefore, I call this view 'Waltonian' perceptualism (see also Ransom, 2020b).

3. Waltonian perceptualism

Waltonian perceptualism unites three main strands that will be elaborated on below. First, perceptual learning explains how aesthetic training allows us to come to perceive artworks in categories, and how appreciation of everyday and natural objects appears to be

untrained. Second, a focus on the way in which perceptual processing takes place – here explored via the fluency hypothesis – explains how learned categories generate the perceptual experiences of (some) aesthetic properties. Third, an account of perceptual expertise provides the resources for determining when such perceptual experiences are veridical.

3.1 *Perceptual learning*

Perceptual learning provides an alternate way of understanding how training is compatible with aesthetic perception (see also Burnston, 2017). Following 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. This change must make a difference to how or whether something is perceived, and it must be brought about as a result of learning, which discounts changes due to lesions and aging, for example.

Some perceptual learning is wholly explicable in terms of low-level properties, such as the improved ability to discriminate the direction of motion of dots on a screen (Ball & Sekuler, 1982). However, there is empirical work that suggests that we can come to represent high-level properties in perceptual experience as the result of perceptual learning (Gauthier et al., 2010), as well as philosophical arguments to this effect (Ransom, 2020a; though see Connolly, 2014, 2019). The same processes involved in the perceptual learning of low-level properties are hypothesized to allow for the representation of high-level properties, such as – I will propose here – categories of art in perceptual experience.

Two such processes of perceptual learning are attentional weighting and stimulus or topographical imprinting (Goldstone, 1998). In attentional weighting, those features or feature dimensions relevant to a given category come to be automatically weighted in attention more heavily over time, and those that are irrelevant come to be weighted less heavily. Features are single properties such as ‘blue’, or ‘square’, whereas dimensions are sets or ranges of features, such as colour, shape, or size. Attentional weight corresponds to how diagnostic a given perceptible feature or dimension is for category membership. So for example, brushstroke style and colour palette are heavily diagnostic of membership in the category of art ‘Impressionist painting’, and thus central to the attribution of the high-level property of *being an Impressionist painting*.¹⁶ Correspondingly, after being shown many exemplars from the category of Impressionist paintings, these features become more heavily weighted in attention, as they are typical of and so diagnostic of category membership. This is similar to how one might learn – after being shown many exemplars of cats – that having whiskers and a long tail are typical perceptible features of this natural kind, useful for categorizing new instances as cats.

Attentional weighting is importantly linked to feature detection: one must be able to detect a given feature in order to accord it more or less weight. In the perceptual learning process of stimulus imprinting, detectors are created for groups of features (or feature dimensions).¹⁷ In topographical imprinting the relations between these features are also

¹⁶ While Walton holds that we can only perceive properties such as ‘in-the-style-of Impressionism’, or ‘apparent etchings’ here I take the stronger stance that we can represent properties such as ‘Impressionist painting’ or ‘etching’ so long as the category possesses distinctive perceptual features. The potential existence of fakes or stylistic copycats should be no more troubling than that of barn facades or decoy ducks.

¹⁷ The perceptual learning process of unitization is another way of grouping feature detectors together. It presumes that the starting point is a bunch of separate feature detectors, which are then unified.

registered in a way that encodes their spatial structure. For example, in the case of cats, not only is it important to detect whiskers and a long tail, it also matters how those features are related to each other – a ‘tail’ with some ‘whiskers’ sprouting out the top would likely disqualify something from being a cat. While the parts are there they do not stand in the right relations to each other. In the case of some categories of art, structural relations may also be important. For example, the elongated necks and stylized faces of the subjects of Modigliani’s portraits are highly characteristic of his individual artistic style, and thus partially diagnostic of this category of art (perhaps, *being a Modigliani portrait*).

The result is a weighted network of feature detectors designed to track the high-level property of interest, whether this be ‘cat’ or ‘Impressionist painting’. This structure should be understood, on the account I am putting forth here, to be equivalent to a prototype (though see Seeley, 2020). Prototypes encode in some way the central tendencies of category members, and use this encoded information in order to categorize objects (Smith & Medin, 1981). While the exact details of what is encoded vary on different accounts of prototypes, in general they are perceptual features or dimensions of objects that are typical or highly diagnostic of a category. In order to be placed into a category, an object must pass a stipulated threshold by possessing a certain number of these features, to some degree. Objects can thus be better or worse examples of cats or Impressionist paintings, while still being categorized as such.

Here readers might wonder exactly how perceptual learning differs from cognitive permeation, especially given that attention has been invoked in both cases. The main difference concerns the role of art-historical knowledge in learning. While possession of art-historical knowledge is necessary for cognitive permeation to occur, it is not required for perceptual learning to occur. For example, in perceptual learning the process of attentional weighting can be accomplished without any top down attentional control through a process known as ‘blind flailing’.

Blind flailing, in its original use, refers to the process by which infants learn to execute skilled movements. At first they randomly flail their limbs and then gradually refine their movements in accordance with whether they result in some sort of reward, such as reaching a toy. In perceptual learning, the term refers to the process whereby attentional weights are initially assigned to different features of an object randomly (Goldstone et al., 2011). Those weightings that are rewarding, insofar as they are useful for tasks such as categorization, are then reinforced and preserved. Those weightings that are not rewarding are revised. This iterative process continues until the weights stabilize, at which point the person has achieved some degree of expertise. In both the perceptual learning and the infant development cases there is a bootstrapping process that does not require the guidance of background knowledge.

This is not to say that art-historical knowledge can play no role in perceptual learning. It may serve as a catalyst to perceptual learning by, for example, guiding our attention to the properties that ought to be more heavily weighted for categorization, and so cut down on the time it takes for us to form a prototype. However, this guided attention does not count as cognitive permeation – it speeds up a process that would have occurred anyways absent such knowledge. While art-historical knowledge may be the cause of our repeated attention to the relevant properties for categorization, the learned attentional weighting is nevertheless set by our exposure to their repeated presence in multiple artworks. We are picking up on statistical properties of the environment, something that can be done without any

background knowledge at all – as corroborated by the abilities of non-human animals and machine learning algorithms discussed in section 2.3.

Moreover, this accelerated pace of learning is not plausibly explained by positing that beliefs permeate perceptual experience. Here I endorse and briefly summarize the arguments of (Ransom, 2020c) in considering what would occur in the case where a person's art-historical beliefs are false. Suppose these beliefs are false in that they guide the person's attention to properties that are not in fact diagnostic of any actual (or potential) category – there is no perceptible pattern to be learned across the set of exemplars. In this case, while such beliefs would guide attention (at least at first), they would not result in the formation of a perceptual category of art, for lack of a perceptible pattern to rely on. We could therefore not reliably come to categorize works of art based on these irrelevant features. In this case, the art-historical beliefs exert no internal causal power on perceptual experience in the form of cognitive permeation. However, both true and false beliefs will affect our actions in the same way, all else equal. If I truly or falsely believe that there is someone at my doorstep, I will open the door. So given that art-historical beliefs exert no causal power in the case of false beliefs, we should not posit that they do so in the case of true beliefs. Rather, this counterfactual scenario supports the claim that the source of the change in the case where one happens to hold true art-historical beliefs is becoming attuned to perceptible features diagnostic of the category (aka perceptual learning).

Perceptual learning thus provides a genuine alternative to cognitive permeation. It explains how training allows us to perceive artworks in categories without relying on art-historical knowledge. While this account does not yet provide an explanation of how we come to perceive aesthetic properties themselves, it nevertheless contains the seed of the solution. Prototypicality is an important factor that determines perceptual processing dynamics, and recent empirical work links such dynamics to our perceptions of beauty and ugliness.

3.2 Perceptual processing dynamics

As per the discussion in section 2.3, proponents of aesthetic perception cannot plausibly hold that we perceive aesthetic properties in the same way that we perceive telephones. There is no general mapping from perceptible non-aesthetic properties to aesthetic properties. This is, in part, what makes exclusive reliance on supervenience so unsatisfying. It is highly plausible that aesthetic properties will depend in some way on low-level perceptible properties but given the lack of clear rules for moving from one to the other, little progress has been made. One alternative is therefore to turn from low-level perceptual content to the *way* in which perceptual processing takes place. The phenomenon of perceptual processing fluency is the most developed and studied in the empirical literature. While it may not ultimately be the correct account of how aesthetic properties are perceived, by focusing on processing dynamics I believe it is on the right track.

The fluency hypothesis, first proposed by Reber, Schwarz and Winkielman (2004), can be broken down into two main parts: (i) an empirical hypothesis that the fluent processing of a perceptual stimulus gives rise to positive affect, and (ii) a metaphysical claim that this

positive affect can be identified with aesthetic pleasure or beauty. In what follows I do not take on board their metaphysical claim, only (i).¹⁸

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 kinds of factors that influence processing fluency.¹⁹

First, some stimuli can be processed more easily than others in virtue of properties of the object that do not require training to detect. These include properties that have been singled out by proponents of objective accounts of beauty as being responsible for our aesthetic attributions. For example, symmetry has long been pointed to as a factor that contributes to beauty, independently of a person's experience with a category (Arnheim, 1954). Humans generally have a preference for symmetric over non-symmetric patterns, all else equal, and in the case of faces, symmetry (or lack thereof) is one of the factors that influences perceived attractiveness or beauty (Rhodes et al., 1999). The fluency explanation is that symmetric stimuli are easier to process because they contain redundant information. 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.²⁰

The second factor that influences processing fluency concerns properties of the object that do require training to detect. This training corresponds to the discussion of perceptual learning in the previous section.²¹ The learning of the characteristic structural relations between features of a stimulus – in perceptual learning what is 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 (Sollberger & Reber, 2004) 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.

There is also some evidence that the more prototypical or standard features that an object possesses relative to a given category, 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 (e.g. Rhodes & Tremewan, 1996).

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

¹⁸ While a full discussion of the metaphysics of aesthetic properties is beyond the scope of this paper, I take it that an objective account of beauty is compatible with the view I set out in this paper. Positive affect may merely be what allows us to detect, and so represent in perception, the property of beauty.

¹⁹ In what follows, my analysis diverges slightly from Reber et. al (2004), as they discuss objective vs. subjective properties of objects rather than trained vs. untrained.

²⁰ For an examination of this principle at work in judgments of Cubist paintings, see (Nicki et al., 1981).

²¹ An anonymous reviewer for this journal questions whether we must learn the pattern qua pattern in order to come to perceive aesthetic properties, suggesting that studies on the mere exposure effect support the view that we do not. The findings of these studies are that merely being exposed to a given object increases our liking of it (Zajonc, 1968), though the story may be complicated when it comes to art (Meskin et al., 2013). However, candidate hypotheses to explain the mere exposure effect tend to invoke processing fluency (Bornstein & D'Agostino, 1994; Newell & Bright, 2001) or the learning of underlying stimulus structure, as hypothesized to occur in perceptual learning (Gordon & Holyoak, 1983). So at least some plausible explanations for the mere exposure effect are consistent with the account on offer here.

causes (or is partially constitutive of) our perceptual experience of the object as having the aesthetic property of beauty or ugliness.

A full treatment of how this account might be extended to other aesthetic properties is beyond the scope of this paper. In part, it will depend on empirical research that has yet to be conducted, linking factors that influence perceptual processing dynamics with attributions of aesthetic properties. Nevertheless, here I offer a brief and rather speculative sketch of how we may represent further aesthetic properties in perceptual experience.²² In the case of processing fluency, my contention is that our understanding of fluency needs to be broken down into the distinctive ways in which fluency can be enhanced or hindered, or by reference to the specific properties of the object that cause the enhancement or hindering. For example, take Walton's claim that contrastandard properties influence which aesthetic properties are perceived. A twitching kinetic painting, he suggests, may be perceived as shocking because it is contrastandard for paintings to possess this property. 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.

Moreover, understanding perceptual categories to have a prototypical structure also gives us a wide range of variables to draw upon. As outlined above, typical features diagnostic of category membership can be understood as varying along a range or dimension. The typical size of an object may in fact be represented as a range of sizes, with a threshold for smallest and largest. We might then begin to understand some other 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), or at the extreme limit for how these different properties are related or proportioned. In the case of cats, for example, kittens may be perceived as cute in virtue of their possessing a very small body in relation to a proportionately large head, and very large, wide-set eyes.

Finally, given that fluency relies on the positive or negative affect generated by perceptual processing in order to explain the perception of aesthetic properties, this opens the door to combining the account with other views of how we come to represent aesthetic properties that invoke other, non-perceptual, sources of affect. Bullot and Reber (2013) extend their account in part by holding that conceptual fluency can also be a source of aesthetic pleasure. To move beyond fluency, accounts of humour often invoke the epistemic emotion of surprise in explanations of aesthetic appreciation (Hurley et al., 2011). The epistemic emotions of curiosity and uncertainty are also sometimes plausibly involved in our aesthetic appreciation of artworks and objects across a broad range of categories, perhaps especially in the case of conceptual art (Goldie & Schellekens, 2009).

Calling upon conceptual fluency or the epistemic emotions might at first seem to count against a perceptualist account, but there are several ways of integrating it. One might simply restrict the claim that we perceive aesthetic properties: while aesthetic properties are sometimes perceived, on other occasions they are represented in experience more broadly (Lopes, 2016). The complexity of our art-appreciative practices suggests that art appreciation likely goes beyond simply perceiving aesthetic properties, but a perceptual account can nevertheless explain how such practices get off the ground in the first place, and how we get

²² For alternative proposals see (Bullot & Reber, 2013; Leder et al., 2004).

drawn into such practices. Or, another way of accommodating non-perceptual affect is to hold that it can in some cases alter perceptual processing (Zadra & Clore, 2011), and thereby have an impact on our aesthetic experiences, changing the aesthetic properties we attribute to people or objects (Matey, 2016).²³

Extensions of the account in one or more of the ways described above helps address what appear to be problematic cases for any account that invokes fluency, such as artworks that are processed in perception in a perceptually disfluent manner but are nevertheless experienced as having positive aesthetic properties. Cases might include jazz music, conceptual art, or twelve-tone music. In such cases, it may be the non-perceptual sources of fluency or affect that explain the outcome.

In sum, while the fluency hypothesis does not likely offer a full explanation of how we come to perceive aesthetic properties, it nevertheless illustrates how focusing on the dynamics of perceptual processing rather than on perceptual content may fill the explanatory gap between perceiving artworks in categories and the perceptual experience of some aesthetic properties. The reason why aesthetic properties such as beauty and ugliness depend in part on which category of art a work is perceived in is because the perceptual category influences the perceptual processing dynamics of the object.

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 (or on any account that appeals to processing dynamics) beauty is not an invariant property of an object, but rather is relative to each subject.

To resolve the threat of relativism, one final element must be added to Waltonian perceptualism. 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 art-historical facts, Walton was looking only for a way to say when our perceptions of artworks as belonging to categories are correct.

By contrast, here there is an additional issue: whether or not our categories are biased, despite being correct. 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 is biased in some way. Perhaps someone's perceptual category of 'Picasso paintings' was formed largely – though not 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.²⁴

²³ Whether or not this affective influence amounts to 'cognitive' permeation is beyond the scope of this paper. See (Matey 2016) for reason to think it might.

²⁴ One might object here that the person is simply perceiving the work in the wrong category, and that given the preponderance of blue period exemplars their learned category is not 'Picasso paintings' but rather 'Picasso blue period paintings'. This would make the issue here identical to the issue Walton was trying to address of how to determine the correct category. While I accept that some examples might be accommodated this way, I think that there will also be cases where one's category is biased but correct. However, I acknowledge there are substantive issues concerning how to determine the correct category of art (Laetz, 2010) and some people might see my project in the next section as trying to solve the same issue as Walton. If so, then I would note that my proposed solution adds to on Walton's account. Thanks to an anonymous reviewer for raising this objection.

In order to address this issue, I develop a sketch of what it is to be an aesthetic perceptual expert that provides a means of separating veridical from illusory perceptual experiences of aesthetic properties.

3.3 *Perceptual expertise*

There are many skills that go into making an expert that are non-perceptual (Cavedon-Taylor, 2017; Montero, 2016). What I am interested in is providing an account of what it means to be an aesthetic perceptual expert with respect to some category of object or domain. 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 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.

Reliable perceptual categorization involves a significantly above-chance ability to categorize the object into the category or categories C to which it actually belongs. Here, I adopt the bulk of Walton's normative thesis: the correct category or categories are determined by art-historical facts (I do not adopt his claim that the correct category is determined in part by which category leads to greater aesthetic appreciation of the object). Reliability is required because to perceptually process an object in the wrong category leads to different fluency effects that will likely result in falsidical aesthetic perceptual experiences.

This 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 wrong category. The children's fairytale of the ugly duckling nicely illustrates this sort of mistake. All the other ducks perceive the 'duckling' to be very ugly, but it turns out that they have miscategorized the bird, which is in fact a swanling. As a member of the proper category, the swanling may be beautiful. The ducks apparently are not aesthetic perceptual experts with respect to the category swanling (or else they would not likely have made the mistake in the first place), and so are not in a position to perceive its true aesthetic properties.²⁵

Even when someone is able to correctly perceptually categorise an object, however, there is still a chance for aesthetic error. Sometimes our perceptual categories themselves may be biased, such as when they are formed based on a skewed sample. This potential for bias is the motivation for the condition that the perceptual category or prototype must approximate the true population average for members of a given category. What this condition amounts to is the requirement that the attentional 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, this colour pattern will be a highly diagnostic feature in this case. It is not as strong a diagnostic for dogs more generally, however, whose colouring varies across a wide range over the different breeds, and so this particular colour pattern should not be as heavily weighted (within the dimension of colour) in this case. The result is that the attentional

²⁵ 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 adult swans possess enough beauty-making properties that don't depend on training so that even those who have never seen a swan before can find them beautiful.

weights track the true population statistics for those features, where the weight corresponds to the average prevalence of the feature in a population.²⁶

Requiring that diagnostic features be weighted according to their actual prevalence allows us to distinguish between veridical and illusory cases of aesthetic perception when the object has been correctly categorized. We might perceptually experience some object to be beautiful due to our fluent processing of it, but if we have only been exposed to a skewed sample of exemplars then our perceptual category 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.

An example of this may be North American beauty pageants, where judges are likely perceptually biased towards Caucasian beauty standards at least in part due to the predominance of Caucasians in the local population.²⁷ 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. In this case, the perceptual category (adult female) is biased because the judges have formed the perceptual category through learning from a biased sample. This biased perceptual category will in turn cause mistaken perceptions and judgments of beauty. Waltonian perceptualism, unlike cognitive permeation, thus provides an account of how aesthetic perception can go wrong, even when the artwork has been accurately categorized.

4. Solving the puzzle of training

The puzzle of training is how aesthetic perception can be both learned and unlearned. The first strand of the puzzle – the challenge from aesthetic training – is to explain how aesthetic perception and training are compatible. Waltonian perceptualism resolves this by understanding training in terms of perceptual learning. Aesthetic training involves significant perceptual training with the goal of the construction of perceptual prototypes for new categories of art, which then alters the processing dynamics of individual category members, such as in the case of processing fluency.²⁸ Moreover, artworks require more explicit training because of their relative absence from our developmental environments. Children (at least in Western culture) are typically exposed to relatively few exemplars of categories of artworks and so do not have the opportunity to perceptually learn these categories (Lopes & Ransom, forthcoming). It thus explains why the sorts of categories that require explicit perceptual training are more likely to be categories of art.

Waltonian perceptualism also provides a solution to the second strand of the puzzle – the challenge from aesthetic asymmetry – which is how aesthetic properties can be perceived without training in some objects but not others. Moreover, it does so in a way that explains the ubiquity of untrained aesthetic experience in our lives. The solution here breaks down

²⁶ 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.

²⁷ Thanks to an anonymous reviewer for this journal, who points out that history will likely play a role in perceptual bias as well. My position is that it does, but only indirectly by altering a person's environment (and not, say, via cognitive permeation). For example, official or unofficial racial segregation, as well as racial bias in advertising, will affect a person's exposure to people of other races. In both cases, perceptual learning will be biased due to this uneven exposure to exemplars and so their aesthetic perceptual experience will be distorted.

²⁸ This account also goes beyond Walton's by providing a more precise framework for the notion of graded standard properties. While Walton notes that there will be graded membership in categories of art (1970, p. 342 fn. 10), here the idea is that some standard properties will count more or less for category membership, depending on their attentional weights. The idea that something is a better or worse example of a given category is central to theories of prototypes.

into three parts. First, it denies that some cases of widespread aesthetic perception are truly untrained. Most of us have already undergone perceptual learning for natural objects such as dogs and horses through the prevalence of these objects in our environments: we are all more or less experts with respect to these categories. While our acquisition of such categories is not trained in the sense of requiring the acquisition of background knowledge, it nevertheless requires perceptual learning. Once prototypes for these categories are formed, exemplars of these categories can also be processed more or less fluently according to their prototypicality (amongst other factors), and so we can perceptually experience these objects as beautiful or ugly as well. Since this learning takes place largely through repeated exposure, this gives rise to the mistaken intuition that our aesthetic appreciation of these objects is untrained. In this way Waltonian perceptualism also provides a unified explanation of aesthetic perception, applicable both to objects that are usually the subject of formal aesthetic instruction and those that we are familiar with through long-term exposure in daily life. Second, given that some fluency-enhancing properties do not rely on experience with a category – such as symmetry or informational redundancy, along with other properties yet to be empirically unearthed – then these properties will not require training. This explanation concords with intuitions that we can sometimes veridically perceive aesthetic properties despite lack of experience with a category of object. And third, some of these ‘perceptions’ of aesthetic properties will be illusory. For example, minimal proficiency with a category – where this is understood as the formation of a prototype that nevertheless does not accurately mirror the population statistics of the category – will result in such perceptual illusions. Or, placing something in the wrong perceptual category may lead to falsidical perceptual experiences.

While I cannot provide a full discussion of the consequences for the epistemology of aesthetic judgments, on the view I have presented as we increase our expertise through perceptual learning our perceptual experiences of aesthetic properties will be more often veridical. However, my account may also signify that aesthetic perceptual illusions are widespread in cases of insufficient training, rendering perceptual experience less reliable a basis for judgments, at least for some.

Finally, turning to philosophy of perception more generally, insofar as aesthetic properties are high-level properties, a complete, plausible account of how they come to be perceived provides additional reason to endorse the rich content view of perceptual experience. A better understanding of how we come to perceive aesthetic properties – in particular with an emphasis on processing dynamics rather than content – may also provide an alternate strategy for theorists attempting to determine whether other high-level properties come to be perceived.

5. Conclusion

In this paper I have argued that Waltonian perceptualism provides a better explanation of aesthetic perception than cognitive permeation. An adequate theory of aesthetic perception should resolve the challenge from aesthetic training: how it is that aesthetic perception is compatible with training. It should also answer the challenge from aesthetic asymmetry: why it is that we seem to require training to perceive aesthetic properties in most artworks but not in nature or everyday objects. While cognitive permeation explains how we come to perceive artworks in categories when training involves the acquisition of art-historical knowledge, it does not explain cases of autodidacts whose training is not guided by

such knowledge, nor cases of everyday aesthetic perception with non-art objects and the natural world. Moreover, the explanation of how we perceive aesthetic properties is only partial, in that it relies on a bare supervenience claim.

In contrast perceptual learning explains how it is we are able to categorize objects, including artworks, in perception. Positing that such learned categories exhibit a prototypical structure sets the stage for understanding how perceptual experience of aesthetic properties emerges, as it provides the key link to the second main strand of the account: the dynamics of perceptual processing, here discussed in terms of processing fluency. While aesthetic properties are not themselves prototypes, they are represented in perceptual experience via the *way* in which aesthetic objects are perceptually processed. An understanding of perceptual expertise in terms of acquiring prototypes that accurately reflect the world gives us the resources to say when aesthetic perceptual experiences are mistaken or veridical.

The picture that emerges is one where we can truly perceive beauty and ugliness – and perhaps other aesthetic properties – but only if the object possesses fluency-enhancing properties that do not depend on training to be perceived, or if we are perceptual experts in a given category. For those who are not experts, apparent instantiations of beauty are a perceptual illusion.²⁹

References

- Alston, W. P. (1983). What's wrong with immediate knowledge? *Synthese*, 55(1), 73–95.
- Arnheim, R. (1954). *Art and visual perception*. University of California Press.
- Ball, K., & Sekuler, R. (1982). A specific and enduring improvement in visual motion discrimination. *Science (New York, N.Y.)*, 218(4573), 697–698.
- Bender, J. W. (1995). General but Defeasible Reasons in Aesthetic Evaluation: The Particularist/Generalist Dispute. *The Journal of Aesthetics and Art Criticism*, 53(4), 379–392.
- Bornstein, R. F., & D'Agostino, P. R. (1994). The Attribution and Discounting of Perceptual Fluency: Preliminary Tests of a Perceptual Fluency/Attributional Model of the Mere Exposure Effect. *Social Cognition*, 12(2), 103–128.
- Brady, E. (1998). Imagination and the Aesthetic Appreciation of Nature. *The Journal of Aesthetics and Art Criticism*, 56(2), 139–147.
- Brogaard, B. (2013). Do we perceive natural kind properties? *Philosophical Studies*, 162(1), 35–42.
- Buchner, A. (1994). Indirect effects of synthetic grammar learning in an identification task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(3), 550–566.
- Bullot, N. J., & Reber, R. (2013). The Artful mind meets art history: toward a psycho-historical framework for the science of art appreciation. *Behavioral and Brain Sciences*, 36(2), 123–137.
- Burnston, D. C. (2017). Is aesthetic experience evidence for cognitive penetration? *New Ideas in Psychology*, 47, 145–156.
- Carlson, A. (1979). Appreciation and the Natural Environment. *The Journal of Aesthetics and*

²⁹ My gratitude to Dominic McIver Lopes for many helpful comments and discussions on numerous iterations of this paper (would that all grad students have a supervisor with such patience and wisdom!), as well as to Murat Aydede, Chris Mole, Evan Thompson, Elisabeth Schellenkens, Kendall Walton, and the audience of the 2018 annual meeting of the American Society for Aesthetics (ASA) in Toronto. Special thanks to the anonymous reviewers for this journal, whose constructive and insightful comments helped me vastly improve the paper. This work was funded in part by a SSHRC Bombardier Doctoral Scholarship and an ASA Dissertation Fellowship. Any views, findings, conclusions, or recommendations expressed in this publication do not necessarily reflect those of the ASA.

- Art Criticism*, 37(3), 267–275.
- Carroll, N. (1993). On Being Moved By Nature: Between Religion and Natural History. In S. Kemal & I. Gaskell (Eds.), *Landscape, Natural Beauty and the Arts* (pp. 244–266). Cambridge University Press.
- Cavedon-Taylor, D. (2017). Reasoned and Unreasoned Judgement: On Inference, Acquaintance and Aesthetic Normativity. *The British Journal of Aesthetics*, 57(1), 1–17.
- Connolly, K. (2014). Perceptual Learning and the Contents of Perception. *Erkenntnis*, 79(6), 1407–1418.
- Connolly, K. (2019). *Perceptual Learning: the flexibility of the senses*. Oxford University Press.
- De Clercq, R. (2002). The Concept of an Aesthetic Property. *Journal of Aesthetics and Art Criticism*, 60(2), 167–176.
- De Clercq, R. (2008). The Structure of Aesthetic Properties. *Philosophy Compass*, 3(5), 894–909.
- Dorsch, F. (2013). Non-Inferentialism about Justification—The Case of Aesthetic Judgements. *The Philosophical Quarterly*, 63(253), 660–682.
- Dutton, D. (1979). Artistic Crimes: the problem of forgery in the arts. *The British Journal of Aesthetics*, 19(4), 302–314.
- Eaton, M. M. (1998). Fact and Fiction in the Aesthetic Appreciation of Nature. *The Journal of Aesthetics and Art Criticism*, 56(2), 149–156.
- Firestone, C., & Scholl, B. J. (2014). “Top-Down” Effects Where None Should Be Found: The El Greco Fallacy in Perception Research. *Psychological Science*, 25(1), 38–46.
- Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral and Brain Sciences*, 39(E229).
- Gauthier, I., Tarr, M. J., & Bub, D. (Eds.). (2010). *Perceptual Expertise: Bridging Brain and Behavior*. Oxford University Press.
- Goldie, P., & Schellekens, E. (2009). *Who’s afraid of conceptual art?* Routledge.
- Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology*, 49(1), 585–612.
- Goldstone, R. L., Landy, D., & Brunel, L. C. (2011). Improving Perception to Make Distant Connections Closer. *Frontiers in Psychology*, 2, 385.
- Gordon, P. C., & Holyoak, K. J. (1983). Implicit learning and generalization of the “mere exposure” effect. *Journal of Personality and Social Psychology*, 45(3), 492–500.
- Hopkins, R. (2005). Aesthetics, Experience, and Discrimination. *The Journal of Aesthetics and Art Criticism*, 63(2), 119–133.
- Horgan, T. (1993). From Supervenience to Superdupervenience: Meeting the Demands of a Material World. *Mind*, 102(408), 555–586.
- Hurley, M. M., Dennett, D. C., & Adams, R. B. (2011). *Inside jokes: Using humor to reverse-engineer the mind*. MIT Press.
- Irvin, S. (2008). The Pervasiveness of the Aesthetic in Ordinary Experience. *The British Journal of Aesthetics*, 48(1), 29–44.
- Iseminger, G. (2004). *The aesthetic function of art*. Cornell University Press.
- Jenkin, Z. (2020). The epistemic role of core cognition. *Philosophical Review*, 129(2), 251–298.
- Laetz, B. (2010). Kendall Walton’s “categories of art”: A critical commentary. In *British Journal of Aesthetics* (Vol. 50, Issue 3, pp. 287–306). Oxford Academic.
- Lamarque, P. (2010). *Work and object: explorations in the metaphysics of art*. Oxford University Press.
- Lecoutre, A., Negrevergne, B., Yger, F., Noh, Y.-K., & Zhang, M.-L. (2017). Recognizing

- Art Style Automatically in painting with deep learning. In *Proceedings of Machine Learning Research* (Vol. 77). PMLR.
- Leder, H., Belke, B., Oeberst, A., & Augustin, D. (2004). A model of aesthetic appreciation and aesthetic judgments. *British Journal of Psychology*, 95(4), 489–508.
- Leopold, A. (1949). *A Sand County Almanac and Sketches Here and There*. Oxford University Press.
- Levinson, J. (2005). Aesthetic Properties. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 78, 211–227.
- Logue, H. (2013). Visual experience of natural kind properties: Is there any fact of the matter? *Philosophical Studies*, 162(1), 1–12.
- Logue, H. (2018). Can We Visually Experience Aesthetic Properties? In A. Bergqvist & R. Cowan (Eds.), *Evaluative Perception* (pp. 42–57). Oxford University Press.
- Lopes, D. M. (2014a). *Beyond Art*. Oxford University Press.
- Lopes, D. M. (2014b). Feckless Reason. In G. Currie, M. Kieran, A. Meskin, & J. Robson (Eds.), *Aesthetics and the Sciences of Mind* (pp. 21–36). Oxford University Press.
- Lopes, D. M. (2016). In the Eye of the Beholder. In J. Dodd (Ed.), *Art, Mind, and Narrative: Themes from the Work of Peter Goldie*. Oxford University Press.
- Lopes, D. M., & Ransom, M. (forthcoming). Perception in Practice. *Review of Philosophy and Psychology*.
- Machery, E. (2015). Cognitive Penetrability: a no-progress report. In J. Zeimbekis & A. Raftopoulos (Eds.), *The cognitive penetrability of perception: new philosophical perspectives* (pp. 59–74).
- Margolis, J. (1998). Farewell to Danto and Goodman. *The British Journal of Aesthetics*, 38(4), 353–374.
- Matey, J. (2016). Good Looking. *Philosophical Issues*, 26(1), 297–313.
- Meskin, A., Phelan, M., Moore, M., & Kieran, M. (2013). Mere Exposure to Bad Art. *The British Journal of Aesthetics*, 53(2), 139–164.
- Montero, B. G. (2016). *Thought in action: Expertise and the conscious mind*. Oxford University Press.
- Nanay, B. (2015). *Aesthetics as philosophy of perception*. OUP Oxford.
- Newell, B. R., & Bright, J. E. H. (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.
- Nicki, R. M., Lee, P. L., & Moss, V. (1981). Ambiguity, cubist works of art, and preference. *Acta Psychologica*, 49(1), 27–41.
- Parsons, G., & Carlson, A. (2004). New Formalism and the Aesthetic Appreciation of Nature. *The Journal of Aesthetics and Art Criticism*, 62(4), 363–376.
- Pylyshyn, Z. W. (1999). Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22(3), 341–365.
- Ramsey, J. L., Langlois, J. H., Hoss, R. A., Rubenstein, A. J., & Griffin, A. M. (2004). Origins of a stereotype: categorization of facial attractiveness by 6-month-old infants. *Developmental Science*, 7(2), 201–211.
- Ransom, M. (2019). Frauds, Posers And Sheep: A Virtue Theoretic Solution To The Acquaintance Debate. *Philosophy and Phenomenological Research*, 98(2), 417–434.
- Ransom, M. (2020a). Attentional Weighting in Perceptual Learning. *Journal of Consciousness Studies*, 27(7–8), 236–248.
- Ransom, M. (2020b). Waltonian Perceptualism. *Journal of Aesthetics and Art Criticism*, 78(1),

66–70.

- Ransom, M. (2020c). Expert Knowledge by Perception. *Philosophy*, 95(3), 309–335.
- Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing Fluency and Aesthetic Pleasure: Is Beauty in the Perceiver's Processing Experience? *Personality and Social Psychology Review*, 8(4), 364–382.
- Rhodes, G., Geddes, K., Jeffery, L., Dziurawiec, S., & Clark, A. (2002). Are Average and Symmetric Faces Attractive to Infants? Discrimination and Looking Preferences. *Perception*, 31(3), 315–321.
- Rhodes, G., Sumich, A., & Byatt, G. (1999). Are Average Facial Configurations Attractive Only Because of Their Symmetry? *Psychological Science*, 10(1), 52–58.
- Rhodes, G., & Tremewan, T. (1996). Averageness, Exaggeration, and Facial Attractiveness. *Psychological Science*, 7(2), 105–110.
- Saito, Y. (2001). Everyday Aesthetics. *Philosophy and Literature*, 25(1), 87–95.
- Samuels, C. A., Butterworth, G., Roberts, T., Graupner, L., & Hole, G. (1994). Facial Aesthetics: Babies Prefer Attractiveness to Symmetry. *Perception*, 23(7), 823–831.
- Seeley, W. P. (2020). *Attentional Engines: A Perceptual Theory of the Arts*. Oxford University Press.
- Seeley, W. P., Buell, C., & Sethi, R. J. (2017). Categories of Art and Computers: A Question of Artistic Style Scientific Workflows-Data Analysis-Reproducibility View project Digital Humanities-Visual Stylometry-Artistic Style View project. *American Society for Aesthetics Newsletter*, 37(3), 9–11.
- Sibley, F. (1959). Aesthetic Concepts. *The Philosophical Review*, 68(4), 421–450.
- Sibley, F. (2001a). Aesthetic and Non-aesthetic. In J. Benson, B. Redfern, & J. Roxbee Cox (Eds.), *Approach to Aesthetics: Collected papers on philosophical aesthetics* (pp. 33–51). Clarendon Press.
- Sibley, F. (2001b). *Approach to Aesthetics: Collected papers on philosophical aesthetics* (J. Benson, B. Redfern, & J. Roxbee Cox (Eds.)). Clarendon Press.
- Sibley, F. (2001c). Particularity, Art, and Evaluation. In J. Benson, B. Redfern, & J. Roxbee Cox (Eds.), *Approach to Aesthetics: Collected papers on philosophical aesthetics* (pp. 88–103). Oxford University Press.
- Siegel, S. (2010). *The Contents of Visual Experience*. Oxford University Press.
- Smith, E. E., & Medin, D. L. (1981). *Categories and Concepts*. Harvard University Press.
- Sollberger, B., & Reber, R. (2004). Artificial grammar learning as a potential mechanism for the acquisition of music preference. *Unpublished Manuscript*.
- Stokes, D. (2013). Cognitive Penetrability of Perception. *Philosophy Compass*, 8(7), 646–663.
- Stokes, D. (2014). Cognitive Penetration and the Perception of Art. *Dialectica*, 68(1), 1–34.
- Stokes, D. (2018). Rich Perceptual Content and Aesthetic Properties. In A. Bergqvist & R. Cowan (Eds.), *Evaluative Perception* (pp. 19–41). Oxford University Press.
- Strahovnik, V. (2004). The riddle of aesthetic principles. *Acta Analytica*, 19(33), 189–208.
- Tormey, A. (1973). Critical judgments. *Theoria*, 39(1–3), 35–49.
- Tucker, J. (1997). Photography as witness, detective, and impostor: visual representation in Victorian science. *Victorian Science in Context*, 378–408.
- Walton, K. L. (1970). Categories of Art. *The Philosophical Review*, 79(3), 334–367.
- Watanabe, S. (2012). Animal Aesthetics from the Perspective of Comparative Cognition. In Watanabe S. & Kuczaj S. (Eds.), *Emotions of Animals and Humans* (pp. 129–162). Springer.
- Wilson, C. (2007). *Let's Talk About Love: A Journey to the End of Taste*. Continuum.

Please refer to the final published article in *Synthese* (2022), vol. 200, issue 2, article 127, DOI: 10.1007/s11229-022-03555-8

Wollheim, R. (1970). *Art and its objects*. Penguin Books.

Zadra, J. R., & Clore, G. L. (2011). Emotion and Perception: The Role of Affective Information. *Wiley Interdisciplinary Reviews. Cognitive Science*, 2(6).

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, 9(2, Pt.2), 1–27.

Zangwill, N. (1998). Aesthetic/sensory dependence. *British Journal of Aesthetics*, 38(1), 66–81.

Zangwill, N. (2001). Formal Natural Beauty. *Proceedings of the Aristotelian Society*, 101(1), 209–224.