

VI—GIST!

TIM BAYNE

A central debate in the philosophy of perception concerns the range of properties that can be represented in perceptual experience. Are the contents of perceptual experience restricted to ‘low-level’ properties such as *location*, *shape* and *texture*, or can ‘high-level’ properties such as *being a tomato*, *being a pine tree* or *being a watch* also be represented in perceptual experience? This paper explores the bearing of *gist perception* on the admissible contents debate, arguing that it provides qualified support for the claim that certain kinds of high-level properties—such as *being a natural scene*—can be perceptually represented.

I

In ‘Perception and Its Objects’, P. F. Strawson puts the following description of visual experience into the mouth of a ‘non-philosophical observer’: ‘I see the red light of the setting sun filtering through the black and thickly clustered branches of the elms; I see the dappled deer grazing in groups on the vivid green grass’ (2011, p. 127). The context in which this passage appears suggests that Strawson intends it to capture not only the *objects* of visual experience but also the *ways* in which those objects are presented in perception. The words that Strawson attributes to his naive observer may sound commonplace, but their implications are of course highly controversial (as Strawson was well aware). Perceptual experience can certainly represent the clustered shapes of the elms and the illumination falling on the grazing deer, but can it also represent the elms as elms and the deer as deer? That is not so clear.

The question of whether properties such as *being an elm* and *being a deer* can be given in perceptual experience is at the heart of the ‘admissible contents’ debate (Hawley and Macpherson 2011). On

one side of this debate are conservatives, who operate with an austere conception of the kinds of properties that can characterize perceptual experience. Although conservatives don't speak with a single voice, typical expressions of the view limit the contents of visual experience to such properties as colour, brightness, form, shape, motion and texture. The central target of Strawson's essay, A. J. Ayer, held a conservative conception of perceptual experience, and the view continues to have plenty of advocates (for example, Brogaard 2013; Dretske 2015; Price 2009; Prinz 2013). On the other side of the debate are liberals such as Strawson, who hold a richer conception of the kinds of properties that can be given in perceptual experience. Although liberals don't speak with a single voice either, the kinds of properties that they take to be perceptually represented include *being a tomato* (Peacocke 1983), *being a pine tree* (Siegel 2006) and *being a watch* (Bayne 2009).

The admissible contents debate has generally been conducted with reference to purely philosophical considerations, with phenomenal contrast arguments playing a particularly important role in the recent literature (Siegel 2010). Such arguments involve the claim that the best explanation for an alleged phenomenal contrast between two perceptual states requires appealing to the perceptual representation of a high-level property. Contrast arguments have proven controversial (Koksvik 2015). Some theorists reject the method in its entirety, claiming that it is 'deeply wrong-headed' and 'epistemically worthless' (Burge 2014, p. 583; see also Block 2014). Other theorists allow that the method could in principle provide evidence in favour of high-level perceptual content, but they remain unmoved by existing contrast arguments.

This paper leaves contrast arguments to one side and focuses instead on the case for liberalism provided by what vision scientists call 'gist perception'. But although my concern here is not with phenomenal contrast arguments, I am concerned with a notion of perceptual content that is phenomenological in nature. Perceptual content, as I understand it here, is a matter of how the world *seems* to the subject (Evans 1982; McGinn 1989). As Siegel has put it, 'nothing counts as a content of experience if it does not reflect the phenomenal character of experience' (2013, p. 850; see also Siegel 2010, pp. 6, 78).

I will describe properties that reflect perceptual phenomenology as being represented in *phenomenal content*. Of course, the notion

of a content ‘reflecting’ the phenomenal character of experience hardly wears its meaning on its sleeve, and in due course we will need to distinguish different senses in which a content might reflect the phenomenology of perception. For now, however, I will proceed with an intuitive understanding of this notion.

II

‘*The Meaning of a Scene*’. In his 1964 film, *The Pawnbroker*, Sydney Lumet inserted a brief scene representing the protagonist’s distant memory.¹ Although the scene is presented for less than a third of a second and is unrelated to the flow of the narrative, the viewer has no difficulty grasping its meaning (Biederman et al. 1983). Lumet was exploiting a phenomenon that psychologists refer to as ‘gist perception’. Although no precise definition of gist perception is available, the notion is often glossed as the amount of information that typical observers can absorb ‘in a single glance’ (Oliva and Torralba 2006).

It is common to distinguish two forms of gist: *scene-based gist*, in which the gisty property is attributed to the perceptual scene as a whole, and *object-based gist*, where the gisty property is attributed to particular objects in the scene. With respect to scene-based gist, typical observers are astonishingly adept at determining whether a presented scene is (say) natural or constructed, whether it is an indoor scene or an outdoor scene, and (for natural scenes) whether it represents (say) a forest, a beach, or a mountain (Greene and Oliva 2009; Henderson and Hollingworth 1999; Oliva 2005; Oliva and Torralba 2006; Potter 1975). With respect to object-based gist, typical observers are astonishingly adept at identifying the high-level properties of presented objects, such as whether they are animals or vehicles (Li et al. 2002; Thorpe et al. 1996). I will focus here on scene-based gist.

Gist raises a number of interesting issues for philosophers of perception. One question concerns whether it is possible for perception to be purely gisty. This idea has sometimes been suggested in connection with the interpretation of Sperling’s experiments on the

¹ Available at <https://www.youtube.com/watch?v=OLtnOGTdLO4>.

reportability of briefly-presented visual stimuli (for discussion see Block 2007; Grush 2007; Stazicker 2011). In these experiments, subjects are unable to reliably report the identities of particular symbols that are presented to them, but they are able to identify them as natural language symbols—or at least, as items that *look* like natural language symbols. The perceptual states here are sometimes described as involving ‘generic perception’, which can be thought of as a kind of object-based gist.

My primary interest here is not with the question of whether purely gisty perception is possible, but with the more fundamental question of how exactly we should understand gist in the first place. Is it a genuinely perceptual phenomenon, or is it a matter of post-perceptual judgement? Do we *perceive* presented scenes as natural, or do we merely *judge* them to be natural on the basis of their visual appearance?

One reason for taking a perceptual view of gist seriously is that the relevant sciences appear to treat gist as perceptual. Research into gist is carried out by vision scientists, it is published in journals devoted to visual perception, and papers on gist bear such titles as ‘High-Level Scene Perception’ (Henderson and Hollingworth 1999) and ‘The Time Course of Abstract Visual Representation’ (Tatler et al. 2003). The opening sentence of a representative paper on gist reads, ‘One remarkable aspect of human visual perception is that we are able to understand the meaning of a complex novel scene very quickly even when the image is blurred ...’ (Oliva and Torralba 2006, p. 23). Of course, the fact that psychologists and neuroscientists treat gist as perceptual hardly *proves* that it ought to be so regarded, but it does—it seems to me—give us good reason to take the view seriously.

Four considerations favour a perceptual treatment of gist, two of which are noted by Fish (2103). The first consideration is temporal: gist is detected *extremely quickly*. Whereas objects must usually be presented for about 150 ms in order to be identified (Evans and Treisman 2005; Fei-Fei et al. 2007; Rayner et al. 2009; Tatler et al. 2003), scene-gist can be accurately detected with exposure times that are as short as 20 ms (Thorpe et al. 1996; Fabre-Thorpe 2001; see also Joubert et al. 2007; Gordon 2004; Potter 1975; Schyns and Oliva 1994). Of even more direct relevance here is the time required to *process* gist information. Here too gist seems to

have an advantage over certain representations that qualify as clearly perceptual, for there is evidence that gist can be recovered within 150 ms of stimulus onset (Thorpe et al. 1996). A study by Macé et al. (2009) found a processing advantage for superordinate categories (animal versus non-animal) over basic categories (dog versus bird) in naturalistic scenes, indicating that ‘in the visual domain, the superordinate level may not constitute an abstraction from basic levels as previously proposed, but rather the rudimentary level at which some coarse object representations can be accessed with early crude processing of visual information’ (Macé et al. 2009, p. e5927).

A second motivation for a perceptual treatment of gist detection is that it requires very little in the way of *focal attention* (Biederman 1972; Wolfe 1998). Li et al. (2002) asked subjects to report as quickly and accurately as possible whether images that were presented for only 27 ms contained an animal (or animals) whilst concurrently performing an attentionally demanding task. Performance on the attentionally demanding task was not significantly different from performance when subjects were presented with just the attentionally demanding task, indicating that object-based gist could be detected in the near absence of attention. Rousselet et al. (2002) provided further evidence that object-based gist is not attentionally demanding, by showing that subjects were as quick and as accurate in discriminating animals from non-animals when two natural images were presented together as they were when only one image was presented. On the basis of these studies, the detection of gist is sometimes said to involve no attention at all. That rather bold claim has been undermined by Cohen et al. (2011), but even they admit that ‘natural-scene perception is so efficient and requires so little attention that the perceptual system must be properly taxed if this attentional cost is to be identified’ (Cohen et al. 2011, pp. 6–7).

A third consideration for treating gist as perceptual concerns its role in *directing and guiding perceptual processing*. As Oliva and Toralba put it, identifying the gist of a scene serves to ‘constrain local feature analysis and enhance object recognition in cluttered natural scenes’ (Oliva and Toralba 2006, p. 23). Information about gist can guide the allocation of attention (Malcolm and Henderson 2009; Neider and Zelinsky 2006; Zelinsky and Schmidt 2009) and the direction of eye-movements to provide a more efficient visual analysis of the scene, so that areas of importance are scanned first and in

more detail (Castelhano and Henderson 2007). Objects that are inconsistent with scene gist are detected more slowly and less accurately than those that are consistent with scene gist (Biederman et al. 1982; Friedman 1979; Henderson, Weeks and Hollingworth 1999), while alterations to a scene that change its gist are more easily detected than those that leave its gist intact (Sampanes et al. 2008). These effects require the perception of the gist, and are not obtained when subjects are merely primed with a verbal description of the scene (Biederman et al. 1983).

A fourth reason for treating gist as perceptual concerns *adaptation*. Adaptation occurs when exposure to a certain property biases the perceptual system away from that property, thus producing characteristic after-effects. A well-known example of adaptation is provided by the waterfall illusion: a stationary object will appear to move upwards following the visual presentation of downwards motion. Although the issue is controversial (Storrs 2015), there is some reason to think that certain gisty properties exhibit adaptation. Greene and Oliva (2010) found evidence of after-effects for four global scene properties: *mean depth* ('the scene takes up kilometres of space' versus 'the scene takes up less than a few metres of space'), *naturalness* ('the scene is a natural environment' versus 'the scene is a man-made, urban environment'), *openness* ('the scene has a clear horizon line, with few obstacles' versus 'the scene is closed, with no discernible horizon line') and *temperature* ('the scene environment depicted is a hot place' versus 'the scene environment depicted is a cold place'). Importantly, these effects were also accompanied by changes in the gist-related judgements that subjects were disposed to make. For example, an image that was near the border between the categories of *forest* and *field* was more likely to be classified as a forest when subjects had previously viewed images that were high in openness (and were thus classified as fields), but that same image was more likely to be classified as a field after adaptation to closed images.²

It would clearly be unwise to regard the perceptual status of gist as settled. There are multiple ways of drawing the contrast between

² Burge (2014) notes out that even when adaptation is associated with high-level perception, it is a further question whether the adaptation derives from the representation of a high-level property rather than the low-level properties on the basis of which that high-level property might be computed. However, Greene and Oliva (2010) provide evidence against a low-level account of their findings.

perception and cognition, and the considerations appealed to here represent only a subset of those that might be regarded as relevant. As Masrour (2011) has noted, any attempt to draw a boundary around the perceptual system will depend in part on views about its function, and that is an issue about which there is debate. However, it seems to me that the case in favour of a perceptual treatment of gist is *prima facie* compelling, and that the burden of proof rests with those who take gist to be a purely post-perceptual phenomenon. With that in mind, let us turn now to consider the bearing of gist perception on the admissible contents debate.

III

Twin Earth, Goldilocks, and the Spatial Envelope. Although Fish eschews the notion of ‘phenomenal content’, he does take gist perception to show that there are ‘no compelling reasons to think that the sensory, presentational component of visual experience must be limited to basic properties’ (2013, p. 54). Is he right to do so? Is gist reflected in the phenomenal character of visual experience, and are gisty properties—that is, the properties that are represented in gist perception—genuinely high-level? Gist provides a case in favour of liberalism only if affirmative answers to *both* of these questions can be justified.

Gisty properties certainly *sound* high-level. Subjects report scene gist with such phrases as ‘it’s an urban scene’, ‘it’s a beach’, ‘it’s a forest’, and they report object-based gist with such phrases as ‘it’s an animal’, ‘it’s a vehicle’, and so on. Indeed, even theoretical treatments of gist employ apparently high-level categories, such as ‘naturalness’ and ‘navigability’ (Greene and Wolf 2011). Although there is no algorithm for deciding when a property is ‘high-level’ as opposed to ‘low-level’, it is fairly clear that *being a beach* and *being an animal* ought to be grouped with canonically high-level properties (such as *being a tomato*) rather than those that are canonically low-level (such as *being square*). Showing that gisty properties can figure in phenomenal content would provide at least a partial vindication of liberalism, even if it failed to motivate the claim that being a tomato, a pine-tree or a watch can be perceptually represented.

But are such properties as *being a beach* really ‘reflected’ in the phenomenology of gist perception? There are a number of reasons for doubt.

We can begin with Twin Earth considerations of the kind that have been advanced against the claim that perceptual content can represent natural kind properties (for example, Brogaard 2013; Price 2009; Prinz 2013). Consider a world in which built environments look the way that natural environments look in the actual world, and vice versa. If you were transported to such a world, natural environments would present you with the gist that you currently associate with built environments, and built environments would present you with the gist that you currently associate with natural environments. The property that is common to the perceptual experience of you and your phenomenal twin is not intuitively regarded as the high-level property of *being a built environment*, but is instead more naturally identified with a low-level property involving the spatial appearance that is distinctive of built environments on earth. Intuitively, when you look at a built environment, your phenomenal twin—at least, your phenomenal twin as far as visual experience is concerned—is the person who, on Twin Earth, is looking at a natural environment, rather than the person who is looking at built environments.³

One might be tempted to respond to this objection by suggesting that it presupposes internalism with respect to phenomenal properties, and that phenomenal internalism is not universally accepted. If an organism’s phenomenal properties constitutively depend on its evolutionary or learning history—as phenomenal externalists claim (Dretske 1995; Lycan 2001)—then we couldn’t assume that you and your molecular duplicate on Twin Earth would be phenomenal twins.

Phenomenal externalism is, of course, a controversial view, and I suspect that few liberals are inclined to accept it. But even if they were to embrace phenomenal externalism, that wouldn’t provide them with a viable response to *this* kind of Twin Earth objection. The reason for this is that the Twin Earth objection just considered made no appeal to the notion of an internal duplicate, but rather

³ Parallel considerations can be used to put pressure on the idea that object-based gist involves high-level properties. For example, consider a world in which animals look the way that vehicles look in the actual world, and vice versa.

appealed to the notion of a *phenomenal* duplicate. Thus, the objection takes no stance on whether phenomenal properties are ‘wide’ or ‘narrow’. What the objection attempts to capture is the idea that *even if* phenomenal properties are wide, they don’t seem to be wide in a way that allows them to represent the high-level properties associated with gist perception. Instead, these high-level properties seem to be screened off by the low-level spatial properties associated with distinct forms of gist. In short, appealing to phenomenal externalism does not appear to offer the liberal a viable response to this kind of Twin Earth objection.

A second objection to the liberal account of gist takes the form of what Dretske (2015) calls the *Goldilocks test*. Dretske introduces this test as a way of telling whether an acknowledged difference in a person’s overall experience is visual as opposed to a difference in what the person ‘thinks, understands, believes, expects or knows’. The test involves a novice (*N*), who is unable to recognize pine trees based on their look, and an expert forester (*E*), who is able to recognize pine trees by sight, and thus—according to liberals about perceptual content such as Siegel (2010)—has visual experiences that represent pine trees. Dretske then imagines that *N* is asked to paint some pine trees, and that *E* is asked to judge the paintings. We can, Dretske says, imagine that *E* has one of the following three reactions:

Too Little. *E* complains that *N* has left something out of his painting: the *pine-tree-ness* of the tree. *E* experiences, he sees, *pine-tree-ness* when he looks at the tree, but he doesn’t see it when he looks at *N*’s painting. (Dretske 2015, p. 166)

Too Much. The forester complains that *N* has put too much in the painting. *N* depicted the tree as having five needles in each cluster with long slender cones, something not characteristic of pine trees in general. *N* did so, of course, because the tree he was painting turned out to be a White Pine tree and White Pine trees have fives needles in each cluster and long slender cones. But in doing this *N* did not, the forester complains, depict the *pine-tree-ness* of the tree, something he sees when he looks at the tree. (Dretske 2015, p. 167)

Just Right. *N*’s painting looks just right to the forester. It depicts the tree exactly as he sees it, exactly as would have painted it. Nothing added, nothing subtracted. (Dretske 2015, p. 167)

Dretske holds that only the third of these reactions is appropriate, and that seems right to me too. But he also holds that option three brings with it an implication that is problematic for liberal treatments of perceptual content:

N always and unintentionally includes in his painting of a pine tree properties (*pine-tree-ness*) he does not see in the tree he is painting. This sounds suspicious. If *N* always and inevitably paints *pine-tree-ness* in his paintings of pine trees by arranging colors, shapes, orientations, and sizes (the properties he does see) the way he does, it begins to sound like *pine-tree-ness* is really just an arrangement, a configuration, of simple properties both *E* and *N* experience when looking at pine trees. If this is so, the property of *pine-tree-ness* that *E* sees is simply an arrangement, a spatial structured array of colored shapes that *N* also experiences when he sees pine trees. He just doesn't know that that arrangement is characteristic of pine trees. (Dretske 2015, p. 168)

If the Goldilocks test undermines the thought that perception can represent *pine-tree-ness*, then surely it also undermines the thought that perception can represent *being a natural scene*. To see how, we need imagine only that Dretske's novice is unable to discriminate natural environments from built ones on the basis of their visual appearance.

The two 'philosophical' arguments against the high-level treatment of gist that we have just considered can be complemented by a third, empirically-based objection, which appeals to the leading computational account of gist perception, Oliva and Torralba's (2001, 2006) 'spatial envelope' model (see also Greene and Oliva 2009). Oliva and Torralba show that a number of global features of scenes—such as their 'naturalness', 'roughness' and 'openness'—can be captured by computational models that quantify only over the spatial features of the scene. Such models capitalize on the fact that, for example, straight horizontal and vertical lines dominate built environments, whereas most natural landscapes have textured zones and undulating contours. Thus, scenes that contain a distribution of textured edges and undulating contours will be classed as having a high degree of naturalness, whereas those that are dominated by vertical and horizontal edges will be regarded as built environments.

Appealing to the spatial envelope model, conservatives might claim that just as perceptual experience doesn't represent pine trees

as such but instead represents only the distinctive spatial gestalt of pine trees, so too gist perception doesn't represent natural environments as such but represents only the distinctive 'spatial envelope' that is associated with such environments. Ordinary observers might lack the vocabulary with which to express this spatial envelope, and thus must rely on high-level descriptions ('It looks like a natural scene'), but—the conservative will claim—the perceptual content that they enjoy is none the less fully captured by the low-level properties implicated in this spatial envelope.

IV

Phenomenal Reflection. What might the liberal say in response to the foregoing considerations? To make progress here we need to return to the notion of phenomenal content and the question of what it is for a property to be reflected by perceptual phenomenology.

The notion of phenomenal content is often understood in terms of *metaphysical supervenience*—that is, it holds across individuals and across possible worlds. On this account, a property is reflected by the phenomenology of a perceptual state just in case its representation is metaphysically necessitated by that phenomenal property. On this view, my phenomenal twins—that is, those individuals who instantiate exactly the same phenomenal properties that I do—will share all (and only) my phenomenal contents. Thus, if each of my phenomenal twins is representing its immediate environment as containing a red square of such-and-such a size, then I too must be representing my environment as containing a red square of such-and-such a size.⁴ Let us call phenomenal content thus defined *strong content*, and the notion of phenomenal reflection that it is intended to capture *strong reflection*.

At least *some* contributors to the admissible contents debate have been concerned with strong content (for example, Brogaard 2013),

⁴ I identify phenomenal content with the kind of content that supervenes on phenomenal character, rather than with the kind of content on which phenomenal character supervenes, precisely because I want phenomenal content to be shared by phenomenal twins. If we were to identify phenomenal content with the kind of content on which phenomenal character supervenes, then phenomenal twins could differ in phenomenal contents, for the fact that representational content fixes phenomenal character doesn't entail that phenomenal character fixes representational content.

and it is not hard to see why one might take this notion to be a legitimate object of interest. For one thing, prominent views in the philosophy of mind are committed to the existence of strong content. Consider those versions of representationalism that *identify* phenomenal properties with representational properties of a certain sort (Tye 1995; Dretske 1995). Although such theories typically focus on the claim that the phenomenal character of perception metaphysically supervenes on its representational properties, in so far as they identify phenomenal properties with representational properties of a certain sort they are also committed to the claim that the intentional content of perception metaphysically supervenes on its phenomenal character. It's not surprising, then, that many treatments of high-level perception have focused on strong content.

But although strong phenomenal content clearly represents a legitimate target of interest, there are good reasons to consider *weaker* notions of phenomenal content, notions that don't require that the representational content of a perceptual state is metaphysically supervenient on the phenomenal character. In fact, there are at least three motivations for such an account of perceptual content.

Firstly, there are accounts of perception that take it to be genuinely representational but deny that there is any strong content in the sense discussed above. Consider, for example, David Papineau's (2014) account of sensory experience, according to which perceptual phenomenology is representational in something akin to the way in which written language is representational. Just as natural language representations carry the contents that they do only contingently, so too Papineau holds that mental representations carry their contents only contingently. Although Papineau's view has no place for strong content, he does not deny that perceptual experience is representational—indeed, he *insists* on it.

A second motivation for a weak notion of phenomenal content derives from Fregean accounts of representationalism (Chalmers 2004; Thompson 2009). Such views hold there is a kind of content that metaphysically supervenes on perceptual phenomenology, but—unlike the better-known Russellian forms of representationalism—the Fregean approach equates this content with modes of presentation rather than with properties as such. On the Fregean view, the strong phenomenal content of a perceptual experience is a kind of condition on extension. So, although Fregeans take contents to be metaphysically supervenient

on perceptual phenomenology, they deny that the representation of particular properties is also metaphysically supervenient on perceptual phenomenology. Fregeans believe in strong phenomenal content if one identifies contents with modes of presentation, but not if one identifies contents with properties, for Fregeans allow that perceptual experiences with the same phenomenal character can represent distinct properties.

A third reason for positing a notion of phenomenal content that is free from commitments of metaphysical supervenience concerns the need to account for the phenomenology of perceptual particularity (Montague 2011; Soteriou 2000). A number of theorists have claimed that perceptual content should be specified in purely *general* terms, and that there is no object-involving perceptual content (for example, Davies 1997; McGinn 1989). There is certainly something to the idea that perceptual content has a purely general character, for there is a sense in which the phenomenal character of one's perceptual experience is unaffected by the numerical identity of the objects of one's experience—all that matters are its qualitative features. Nonetheless, there is also something to the idea that perceptual experience has *singular* content, and that specifying the content of my perceptual experience requires referring to the individual objects at which I am looking (if any). Consider a situation in which I am on earth and am looking at Barack Obama while my twin on Twin Earth is looking at Barack Obama's identical twin. Although there is a sense in which there is no phenomenal difference between our experiences, we might also want to say that there is a sense in which our respective experiences reflect *different* objects: my experience 'reflects' Barack Obama, whereas my twin's experience 'reflects' Barack Obama's twin. In order to capture that thought we need a notion of phenomenal content that isn't metaphysically necessitated by the phenomenal properties that we instantiate.⁵

Taken together, these considerations motivate the need for a sense in which a property can be 'reflected' in the phenomenology of perceptual experience without being metaphysically necessitated by it—in other words, we need a notion of *weak content* (see also Davies 1997). What, precisely, it takes for a property to be weakly reflected by the phenomenology of perception is of course a (very) good question. We

⁵ For further discussion of the parallels between the debate about whether perception is singular or general and the debate about whether or not it has high-level content, see Speaks (2015).

might want to appeal to the laws of nature that obtain in the organism's environment, to the organism's evolutionary or learning history, or to features of its current environment (such as the identity of the objects with which it is in perceptual contact). Identifying which of these properties we should appeal to (and when) is an issue that I leave to one side. My concern here is not with what grounds weak content, but with the mere fact that there is good reason to posit such a notion.⁶

V

With the notion of weak content in hand, we can now return to the objections against the liberal treatment of gist. Consider first the Twin Earth objection, involving a world in which natural environments appear to us in the way that built environments on earth do, and vice versa. This objection appears to show that *being a natural environment* cannot be reflected in the phenomenology of visual experience, but we can now see that this objection presupposes a strong reading of phenomenal reflection. If my phenomenal twin and I can differ in the high-level properties that are represented in our perceptual phenomenology, then such properties are clearly not metaphysically supervenient on the phenomenal character of perception. However, it doesn't follow that *being a natural environment* isn't in some sense reflected in the perceptual experiences that *I* have on earth. I am perceptually sensitive to certain spatial envelopes precisely because of the high-level information they provide about my environment, and in this sense these perceptual experiences can be said to reflect those high-level properties. In worlds in which those spatial envelopes are indicative of distinct high-level properties, then perceptual experiences will weakly reflect those properties. These input-based considerations for thinking of gist in high-level terms can be reinforced by considering the role of gist perception in

⁶ Not only is the notion of weak content well motivated in its own right, it also enables us to make sense certain discussions of high-level perceptual experience that would otherwise be puzzling. Although Siegel holds that natural kind properties such as being a pine tree are reflected in the phenomenal character of perceptual experience, she allows that two natural kinds can be represented by the same type of phenomenal state (Siegel 2013, p. 851). This point would be puzzling if Siegel were thinking of *being a pine tree* as part of the strong content of perception, but it is perfectly understandable if she takes such properties to be components of the *weak* content of perception.

guiding (say) eye movements and the allocation of attention. For example, the gist that is associated with natural environments on earth will generate very different patterns of cognitive and behavioural responses in the inhabitants of an earth in which natural environments have the visual appearance that built environments have on earth.

What about Dretske's Goldilocks objection? At the heart of this objection is the idea that the novice could not represent the *pine-tree-ness* of pine trees unless he himself were able to recognize pine trees. Why might one make such an assumption? I suspect that Dretske is thinking along the following lines. Suppose that the novice was completely colour-blind, and as such was unable to perceive the colour of the pine trees that he is painting. In such a case, it would be puzzling if his painting were able to prompt the expert to accurately perceive the colour of the trees. This case seems to provide intuitive support for Dretske's assumption that the perceptual experience elicited by looking at a representation of an object cannot contain content that the author of that representation is not perceptually sensitive to.

But now consider another variant of Dretske's scenario. Suppose that Dretske's novice suffers from an inability to visually perceive depth, and can represent only the two-dimensional projection of a scene. It is this projection, let us suppose, that the novice attempts to represent in his painting. Could the expert nonetheless perceive depth in the novice's painting? If we adopt the Goldilocks test then we would be forced to answer this question in the negative, for the Goldilocks test assumes that one cannot see more in an image than its architect is perceptually sensitive to. But it would surely be a mistake to suppose that the expert couldn't form an accurate representation of the relative size of the represented trees on the basis of looking at the novice's painting. (Berkeley might have held that depth cannot be visually represented, but few contemporary conservatives would want to follow his lead.) Even though the representation of a distant tree occupies less canvas than is occupied by the representation of the more proximal tree, the laws of perspective surely enable the expert to see the distinct tree as larger than the closer tree. It seems to me that the Goldilocks test ought to be rejected.

What, finally, of the spatial envelope objection? The objection shows that the high-level content of gist perception isn't a *primitive* feature of its content. We are sensitive to the gist of a scene in virtue

of being sensitive to its spatial properties, and it is this feature of gist that the spatial envelope model attempts to capture. But the objection goes wrong in assuming that the content of gist is *exhausted* by the features that we appeal to in characterizing a scene's spatial envelope. Indeed, from an explanatory perspective, limiting ourselves to the content of the spatial envelope is to miss the *entire* point of gist perception. We are perceptually sensitive to scenes that are dominated by vertical and horizontal lines precisely because they are indicative of built environments.

VI

In his commentary on Siegel's *The Contents of Visual Experience*, Prinz remarks that although her view aligns with common sense 'insofar as we take ourselves to see lions and tigers and bears ... it breaks from the kinds of information-processing stories that dominate in vision science' (Prinz 2013, p. 827). The burden of this paper has been to suggest that the story told by vision science is not quite as one-sided as Prinz's remarks suggest. Vision science might not support the claim that we perceptually experience lions, tigers and bears, but it does—I have argued—provide some support for the claim that we perceptually represent beaches, forests and cityscapes. The more general lesson to be drawn from the foregoing is that we need to distinguish different senses in which a property can be given in the contents of visual experience, and thus different versions of 'the' admissible contents debate.⁷

Philosophy Department
School of Philosophical, Historical and International Studies
Monash University
Clayton, VIC 3800
Australia

⁷ I am grateful to the members of the audience at the meeting of the Aristotelian Society—in particular, Sam Clarke, Alex Geddes, Tomis Kapitan and Rory Madden—for their questions. Thanks also to Ned Block, Jakob Hohwy, Eric Mandelbaum, Angela Mendelovici, Tom McClelland and Ryan Oglive for very helpful comments on an earlier version of this paper. This paper was supported by a European Research Council grant *The Architecture of Consciousness* (313552).

REFERENCES

- Bayne, Tim 2009: 'Perception and the Reach of Phenomenal Content'. *Philosophical Quarterly*, 59, pp. 385–404.
- Biederman, Irving 1972: 'Perceiving Real-World Scenes'. *Science*, 177, pp. 77–80.
- Robert J. Mezzanotte, and Jan C. Rabinowitz 1982: 'Scene Perception: Detecting and Judging Objects Undergoing Relational Violations'. *Cognitive Psychology*, 14, pp. 143–77.
- Richard C. Teitelbaum, and Robert J. Mezzanotte 1983: 'Scene Perception: A Failure to Find a Benefit from Prior Expectancy or Familiarity'. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(3), pp. 411–29.
- Block, Ned 2007: 'Consciousness, Accessibility and the Mesh between Psychology and Neuroscience'. *Behavioral and Brain Sciences*, 30(5–6), pp. 481–99.
- 2014: 'Seeing-As in the Light of Vision Science'. *Philosophy and Phenomenological Research*, 89(3), pp. 560–73.
- Brogaard, Berit 2013: 'Do We Perceive Natural Kind Properties?' *Philosophical Studies*, 162, pp. 35–42.
- Burge, Tyler 2014: 'Reply to Block: Adaptation and the Upper Border of Perception'. *Philosophy and Phenomenological Research*, 89(3), pp. 573–83.
- Castelhano, Monica S., and John M. Henderson 2007: 'Initial Scene Representations Facilitate Eye Movement Guidance in Visual Search'. *Journal of Experimental Psychology: Human Perception and Performance*, 33, pp. 753–63.
- Chalmers, David J. 2004: 'The Representational Character of Experience'. In Brian Leiter (ed.), *The Future for Philosophy*, pp. 153–81. Oxford: Oxford University Press.
- Cohen, Michael A., George A. Alvarez, and Ken Nakayama 2011: 'Natural-Scene Perception Requires Attention'. *Psychological Science*, 22(9), pp. 1165–72.
- Davies, Martin 1997: 'Externalism and Experience'. In Ned Block, Owen Flanagan and Güven Güzeldere (eds.), *The Nature of Consciousness: Philosophical Debates*, pp. 309–27. Cambridge, MA: MIT Press.
- Dretske, Fred 1995: *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- 2015: 'Perception versus Conception: The Goldilocks Test'. In John Zeimbekis and Athanassios Raftopoulos (eds.), *The Cognitive Penetrability of Perception: New Philosophical Perspectives*, pp. 163–73. Oxford: Oxford University Press.
- Evans, Gareth 1982: *The Varieties of Reference*. Oxford: Clarendon Press.

- Evans, Karla K., and Anne Treisman 2005: 'Perception of Objects in Natural Scenes: Is it Really Attention Free?'. *Journal of Experimental Psychology: Human Perception and Performance*, 31(6), pp. 1476–92.
- Fabre-Thorpe, Michèle, Arnaud Delorme, Catherine Marlot, and Simon Thorpe 2001: 'A Limit to the Speed of Processing in Ultra-Rapid Visual Categorization of Novel Natural Scenes'. *Journal of Cognitive Neuroscience*, 13(2), pp. 1–10.
- Fei-Fei, Li, Asha Iyer, Christof Koch, and Pietro Perona 2007: 'What Do We Perceive in a Glance of a Real-World Scene?' *Journal of Vision*, 7, pp. 1–29.
- Fish, William 2013: 'High-Level Properties and Visual Experience'. *Philosophical Studies*, 162, pp. 43–55.
- Friedman, Alinda 1979: 'Framing Pictures: The Role of Knowledge in Automatized Encoding and Memory for Gist'. *Journal of Experimental Psychology: General*, 108, pp. 316–55.
- Gordon, Robert D. 2004: 'Attentional Allocation During the Perception of Scenes'. *Journal of Experimental Psychology: Human Perception and Performance*, 30, pp. 760–77.
- Greene, Michelle R., and Aude Oliva 2009: 'Recognition of Natural Scenes from Global Properties: Seeing the Forest without Representing the Trees'. *Cognitive Psychology*, 58, pp. 137–76.
- 2010: 'High-Level Aftereffects to Global Scene Properties'. *Journal of Experimental Psychology: Human Perception and Performance*, 36, pp. 1430–42.
- Greene, Michelle R., and Jeremy M. Wolfe 2011: 'Global Image Properties Do Not Guide Visual Search'. *Journal of Vision*, 11(6), pp. 1–9.
- Grush, Rick 2007: 'A Plug for Generic Phenomenology'. *Behavioral and Brain Sciences*, 30(5–6), pp. 504–5.
- Hawley, Katherine, and Fiona Macpherson (eds.) 2011: *The Admissible Contents of Experience*. Oxford and Malden, MA: Wiley-Blackwell.
- Henderson, John M., and Andrew Hollingworth 1999: 'High-Level Scene Perception'. *Annual Review of Psychology*, 50, pp. 243–71.
- Henderson, John M., Phillip A. Weeks, Jr., and Andrew Hollingworth 1999: 'The Effects of Semantic Consistency on Eye Movements During Complex Scene Viewing'. *Journal of Experimental Psychology: Human Perception and Performance*, 25, pp. 210–28.
- Joubert, Olivier R., Guillaume A. Rousselet, Denis Fize, and Michèle Fabre-Thorpe 2007: 'Processing Scene Context: Fast Categorization and Object Interference'. *Vision Research*, 47, pp. 3286–97.
- Koksvik, Ole 2015: 'Phenomenal Contrast: A Critique'. *American Philosophical Quarterly*, 52(4), pp. 321–34.

- Li, Fei Fei, Rufin VanRullin, Christof Koch, and Pietro Perona 2002: 'Rapid Natural Scene Categorization in the Near Absence of Attention'. *Proceedings of the National Academy of Sciences (USA)*, 99(14), pp. 9596–601.
- Lycan, William G. 2001: 'The Case for Phenomenal Externalism'. *Philosophical Perspectives*, 15: *Metaphysics*, pp. 17–35.
- Macé, Marc J.-M., Olivier R. Joubert, Jean-Luc Nespoulous, and Michèle Fabre-Thorpe 2009: 'The Time-Course of Visual Categorizations: You Spot the Animal Faster than the Bird'. *PLoS ONE*, 4(6), e5927. doi:10.1371/journal.pone.0005927.
- Malcolm, George L., and John M. Henderson 2009: 'The Effects of Target Template Specificity on Visual Search in Real-World Scenes: Evidence from Eye Movements'. *Journal of Vision*, 9(11), pp. 1–13.
- Masrour, Farid 2011: 'Is Perceptual Phenomenology Thin?' *Philosophy and Phenomenological Research*, 83, pp. 366–97.
- McGinn, Colin 1989: *Mental Content*. Oxford: Blackwell.
- Montague, Michelle 2011: 'The Phenomenology of Particularity'. In Tim Bayne and Michelle Montague (eds.), *Cognitive Phenomenology*, pp. 121–40. Oxford: Oxford University Press.
- Neider, Mark B., and Gregory J. Zelinsky 2006: 'Scene Context Guides Eye Movements During Search'. *Vision Research*, 46, pp. 614–21.
- Oliva, Aude 2005: 'Gist of the Scene'. In Laurent Itti, Geraint Rees and John K. Tsotsos (eds.), *Neurobiology of Attention*, pp. 251–6. San Diego, CA: Elsevier.
- 2006: 'Building the Gist of the Scene: The Role of Global Image Features in Recognition'. *Progress in Brain Research*, 155, pp. 23–36.
- and Antonio Torralba 2001: 'Modeling the Shape of the Scene: A Holistic Representation of the Spatial Envelope'. *International Journal of Computer Vision*, 42(3), pp. 145–75.
- Papineau, David 2014: 'Sensory Experience and Representational Properties'. *Proceedings of the Aristotelian Society*, 114(1), pp. 1–33.
- Peacocke, Christopher 1983: *Sense and Content*. Oxford: Oxford University Press.
- Potter, Mary C. 1975: 'Meaning in Visual Scenes'. *Science*, 187, pp. 965–6.
- Price, Richard 2009: 'Aspect-Switching and Visual Phenomenal Character'. *Philosophical Quarterly*, 59, pp. 508–18.
- Prinz, Jesse 2013: 'Siegel's Get Rich Quick Scheme'. *Philosophical Studies*, 163, pp. 827–35.
- Rayner, Keith, Tim J. Smith, George L. Malcolm, and John M. Henderson 2009: 'Eye Movements and Visual Encoding During Scene Perception'. *Psychological Science*, 20, pp. 6–10.

- Rousselet, Guillaume A., Michèle Fabre-Thorpe, and Simon J. Thorpe 2002: 'Parallel Processing in High-Level Categorization of Natural Images'. *Nature Neuroscience*, 5(7), pp. 629–30.
- Sampanes, Anthony Chad, Philip Tseng, and Bruce Bridgeman 2008: 'The Role of Gist in Scene Recognition'. *Vision Research*, 48, pp. 2275–83.
- Schyns, Philippe G., and Aude Oliva 1994: 'From Blobs to Boundary Edges: Evidence for Time- and Spatial-Scale-Dependent Scene Recognition'. *Psychological Science*, 5, pp. 195–200.
- Siegel, Susanna 2006: 'Which Properties are Represented in Perception?' in Tamar Szabó Gendler and John Hawthorne (eds.), *Perceptual Experience*, pp. 481–503. Oxford: Oxford University Press.
- 2010: *The Contents of Visual Experience*. Oxford University Press.
- 2013: 'Replies to Campbell, Prinz, and Travis'. *Philosophical Studies*, 163, pp. 847–65.
- Soteriou, Matthew 2000: 'The Particularity of Perceptual Experience'. *European Journal of Philosophy*, 8(2), pp. 173–89.
- Speaks, Jeff 2015: *The Phenomenal and the Representational*. Oxford: Oxford University Press.
- Stazicker, James 2011: 'Attention, Visual Consciousness and Indeterminacy'. *Mind and Language*, 26(2), pp. 156–84.
- Storrs, Katherine R. 2015: 'Are High-Level Aftereffects Perceptual?' *Frontiers in Psychology*, 6(157). doi: 10.3389/fpsyg.2015.00157.
- Strawson, P. F. 2011: 'Perception and its Objects'. Reprinted in his *Philosophical Writings*, edited by Galen Strawson and Michelle Montague. Oxford: Oxford University Press.
- Tatler, Benjamin, Iain Gilchrist, and Jenny Rusted 2003: 'The Time Course of Abstract Visual Representation'. *Perception*, 32, pp. 579–93.
- Thompson, Brad 2009: 'Senses for Senses'. *Australasian Journal of Philosophy*, 87(1), pp. 99–117.
- Thorpe, Simon, Denis Fize, and Catherine Marlot 1996: 'Speed of Processing in the Human Visual System'. *Nature*, 381, pp. 520–2.
- Tye, Michael 1995: *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- Wolfe, Jeremy M. 1998: 'What Do You Know About What You Saw?' *Current Biology*, 8, pp. R303–4.
- Zelinsky, Gregory J., and Joseph Schmidt 2009: 'An Effect of Referential Scene Constraint on Search Implies Scene Segmentation'. *Visual Cognition*, 17, pp. 1004–28.