

Loar on lemons: the particularity of perception and singular perceptual content

This paper advances two claims. The positive claim offers a correctness condition for perceptual experiences, one that does justice to the so-called “particularity of perception”:

- (T1) the perceptual content of a perceptual experience is correct iff there are perceived objects of which it is non-accidentally true.

The negative claim is that, subject to a qualification entered later (§7):

- (T2) There is no singular perceptual content,

where singular content can be defined as follows:

- (T3) A representational vehicle has singular content iff there is a unique object it rigidly designates.¹

The claims are connected by the idea that (T1) shows that we *do not need* to posit singular perceptual content in order to do justice to the particularity of perception, and (T2) shows that in any case we *could not* use singular content in that way.

I shall assume that to perceive is to represent, and that we perceive, track, and count individual objects. These premises do not entail that there are perceptual representational vehicles with singular content. The distinction between qualitative and numerical identity is available conceptually, but not perceptually. Perceptual representation is not sensitive to numerical identity, and this blocks it from having singular representational content. I reach this conclusion in the light of Brian Loar’s rich and suggestive discussion in “Phenomenal intentionality as the basis of mental content”.²

I’ll begin by showing, in §1, how the issues to be discussed have been distorted by the implicit assumption that there are just two kinds of content: general content, corresponding conceptually to quantified content, and singular content, corresponding conceptually to the content possessed by singular demonstrative concepts or proper nominal concepts. By

¹ This could be nuanced in at least two ways. One might wish to expand it to allow for empty cases, ones in which the contents count as singular in virtue of “purporting” to rigidly designate, even though they in fact designate nothing. And one might wish to narrow it to apply only to simple representational vehicles.

² Loar 2003, though he does not draw the same conclusion.

contrast, I'll suggest that there is at least one other kind of content, alluded to in (T1) above, which I call "feature-expressing". Although I will have to use linguistic predicates in order to point towards feature-expressing contents, this should not be taken as an endorsement of the view that perceptual representational vehicles are predicate-like, or indeed in any way language-like. In (§2), I'll show that measuring instruments are appropriately associated with feature-expressing correctness conditions on the lines of (T1). This will give those who see value in the comparison between perceptual systems and measuring instruments a reason to accept (T1)³. In the remainder of the paper I offer direct arguments for (T1) and (T2), starting in §3 by recalling Loar's example of the shifting lemons. Subsequent sections amplify some related theoretical issues, and the final section offers a qualification of some earlier claims, and has something to say about the nature of feature-expressing contents.

1. Feature-expressing content

Perceptual experiences have representational content. I take this widely-held thesis for granted.⁴ More controversial is the nature of this content. We experience particular things, trees, people, chairs. Representationalists about perception may well suppose that we have perceptual representational vehicles with singular content: their nature is to rigidly designate just one individual object and not anything else. Not only is this a natural view, there are arguments which apparently make it compulsory. Salient among these are mirror cases, which go back to Grice (1961); similar examples have been sympathetically discussed more recently by Soteriou (2000: 180) and Tye (2018) among others. In Grice's example, I am standing in front of one of several pillars. An angled mirror obscures the pillar straight ahead of me, reflecting the image of an indistinguishable pillar off to one side. Despite the correspondence between the qualitative aspects of my experience and a pillar ahead, it is intuitive that some kind of error has occurred. In some sense, I undergo misrepresentation: my perceptual content is inaccurate.

³ "perceptual systems are thus brought under the umbrella of measuring instruments generally (which is surely where they belong ...)" (Churchland 2013: 90).

⁴ The paper assumes only a version of representationalism that, according to Pautz, "is compatible with nearly every view on experience", according to which experiences "merely *have* contents in some sense", leaving open whether they are "*identical with* relations to contents" (2010: 254).

The example shows that the content of my perception cannot be purely general, for example expressible by “There’s a pillar ahead”, since this content is true. A frequently made suggestion is that there is singularity in the content. It is more like “That pillar is ahead” (as in Burge 1991 and McDowell 1991: 266–7). We cannot understand “that pillar” as referring to the one in front of me, since I am visually isolated from it. If the content corresponding to “that pillar” is true of anything, it is true of the reflected pillar off to the side, rather than ahead, so the total content is not true, so my experience is mistaken; and this intuitively is the right result.

Arguments against existentially general content only favor singular content if these are the only alternatives. But they are not. Examples like the pillar show beyond doubt that the content is not merely existentially quantified. But this does not establish that the content is singular, for there are forms of content that are neither existentially quantified nor singular. One example of such content is the feature-expressing content that figures in (T1). These contents can properly be said to be, or fail to be, true *of* objects, and they are linguistically expressed by predicates or open sentences. They are properties or features, and taken by themselves are not evaluable as true or false. According to (T1), perceptual representations that express features as contents are central to perceiving. The work supposedly done by singular perceptual content is instead done by the relation between a perceiver and something perceived.⁵

The correctness condition stated by (T1) is, in a sense, object-involving: a perceptual experience is correct iff there are objects of which its content is correct. This might seem to lead right away to the conclusion that some perceptual content is singular, but the transition would be confused: (T1) makes no mention of any perceptual representational vehicles with singular content. According to (T1), the particularity of perception is ensured by perceiving particulars, not by deploying singular contents.

(T1) provides a starting point for the usual distinctions between veridical, illusory and hallucinatory experience. A veridical experience has perceived satisfiers of its content; an

⁵ Some motivations for singular content are satisfied by (T1). For example, Soteriou writes “When a subject perceives an object, the veridicality of the experience is dependent on how things are with the object being perceived” (2000: 184). This is (approximately) entailed by (T1). But Soteriou also goes further than this premise: “The particularity of a successful visual experience should be reflected in the intentional content of the experience” (2000: 175).

illusory experience has perceived objects which fail to satisfy the content; a hallucinatory experience involves no perceived objects.

Applied to the pillar case, (T1) gives the intuitively correct result. The pillar ahead is not perceived, so it cannot enter the correctness condition stated by (T1). The pillar to one side is perceived (at least, on many views), but it does not satisfy the feature-expressing perceptual content *is ahead*. (T1) correctly rules the perceptual experience to be mistaken, but does not attribute singular content.

2. Measuring instruments

A car's gas gauge represents the level of gas in the tank.⁶ It represents correctly just if the fuel in the connected tank is at the level represented. The gauge contains no proprietary representation of the tank itself, no representational vehicle whose reference is the tank. The needle does not specifically represent the tank, nor does its position on the dial do so, nor does the dial or its markings. The very same gauge, with its needle in the same position, would represent the level of the gas in a different tank were it properly connected to that tank rather than to the tank it is actually connected to. If the needle points to a position marked "1/2", we should view the gauge's proprietary content as something like *is half full*, correct iff true of the connected tank, that is, iff that tank is half full. A gas gauge not connected to a tank does not have a correct content (and perhaps not an incorrect one either). In summary:

(T4) A measuring instrument measures correctly, on an occasion, iff there is suitably causally related object of which the measurement is non-accidentally true.

Just as the gauge contains no representational vehicle that refers to the connected tank, so perceptual experience contains no representational vehicle that refers to perceived objects. This does not prevent a state of the gauge from being evaluated as correct or incorrect relative to a properly connected gas tank; likewise, it does not prevent a perceptual state from being evaluated as correct or incorrect relative to perceived objects. In both cases, an object involved in the correctness condition is one related by a suitable causal and information-carrying link to the representational system.

⁶ My discussion is influenced by Dretske's writing, e.g. his 1995.

A feature-expressing content like *is half full* lacks truth conditions, and so is neither true nor false. We can make sentential attributions of content to measuring instruments, saying for example that the gauge represents *that* the tank is half full. The attributed content is true or false, depending as the connected tank is or is not half full. Different elements of the attributed truth conditional content arise in different ways. The position of the needle is a representational *vehicle*; how the gauge is calibrated determines its representational *content*. Some gauges are calibrated so that having the needle in the vertical position means *full*, others so that the same position means *half full*. The gauge has no tank-specific representational vehicle; the tank enters the correctness condition by being suitably connected to the gauge. The gauge has no internal feature specific to that gas tank. The truth conditional content is hybrid, determined in part by a content for which the gauge has a specific vehicle, and in part by the gauge's relation to its surroundings. The separability of these elements of the hybrid is shown by what an intrinsically unchanged gauge would represent if connected to another tank.

The thermometer reads 85°. If its tip is in your coffee, it represents the temperature of your coffee. But it has no representational vehicle related to your coffee. It knows nothing of you or of coffee. Its representational vehicles are the positions of the indicator liquid and the associated markings. If the top of the liquid is alongside "85", and the calibration is for Fahrenheit, the content is *is 85° F*. If the thermometer is correctly calibrated, that is, if its vehicles have contents that vary systematically with temperatures surrounding the tip, we can infer that the region surrounding the tip is at 85° F. The correctness condition does not require the content to contain an element dedicated to representing that region. Moving the thermometer elsewhere does not in itself change its specific representational contents. There is also a truth conditional attribution to be had: in the envisaged circumstances, the thermometer represents *that* your coffee is at 85° F. This propositional representation is achieved by the union of dissimilar means. The temperature is explicitly and internally represented by the level of the liquid. Its attribution to your coffee comes from a causal fact external to the thermometer itself: its tip is in your coffee.

If one of your perceptual states involves the content *is a green apple*, then intuitively, and by (T1), it is correct just if you are perceiving a green apple. There may also be sentential

attributions to be had, for example: you see that the apple before you is green. This attributes a true or false content. But, as in the case of instruments, the content is hybrid. The *is a green apple* element is an intrinsic feature of the perceptual state, a content for which the state has a specific vehicle. Correctness is measured by whether this content is true of something perceived. This view of perception is encouraged by the analogs with instruments. Later I provide a direct argument for it.

3 Loar's lemons

Content is a theorist's notion, and its use must be justified by the role it plays in a theoretical account of the subject matter. Alleged intuitions about content stand in need of justification in terms of theoretical utility. If it is not theoretically useful to ascribe a content to something, we should not regard that thing as possessing content. In the case of perception, one would expect singular perceptual content, if there is such a thing, to make a difference either to phenomenology or to causal role. But in fact it is hard to see how it could make a difference to either.

Many people (including myself) would be happy to accept the following principle connecting phenomenal character and content:

- (T5) Experiences with subjectively the same phenomenal character have the same representational content.

There is a straightforward argument from (T5) to (T2), the claim that there is no singular perceptual content. One experience might be subjectively indistinguishable from another, even if a perceived object in one is replaced by an indistinguishable perceived object in the other. This indistinguishability suggests identity of phenomenal character (or so many will assume—but see footnote 8 below). If there were singular perceptual content, it would differ in such cases, involving one of the indistinguishable objects in one case, the other in the other; and this is inconsistent with (T5).

However, (T5) is controversial, and so I cannot rely upon it. Indeed, the argumentative direction might be reversed: subjectively indistinguishable experiences intuitively have different correctness conditions if they involve different objects. It might seem that only by appeal to

different singular contents can this fact be accommodated. And this would be inconsistent with (T5).

Another relevant principle connects causal role and content, and many theorists would find it appealing. This version is owed to Neander (2017: 192):

(T6) Differences in sensory-perceptual contents must be accompanied by differences in the internal causal roles of the relevant representations.

Perceptions of duplicates at different moments would not differ in their narrow causal role, and so (T6) would rule that they do not differ in content, which in turn means that these are not cases of singular perceptual content.

The disjunction of (T5) and (T6) makes a powerful package. Both exclude singular perceptual content, but if such content makes no difference either to phenomenology or to narrow functional role, serious doubt is cast on its explanatory value. Given the theoretical nature of content, to say that singular perceptual content has no explanatory value is to say that there is no such thing.

The remaining possibility is that singular perceptual content makes a difference to wide functional role. This might enable one to register in perception a specific member of one's own family or a specific location, and these would be valuable abilities. But (T1) shows that an explanation of them in terms of content is unnecessary: instead, an explanation based on which objects are perceived works fine. And the purely qualitative nature of the senses makes it impossible for them to register genuinely singular content, for they cannot register distinct duplicates in distinct ways.

These abstract reflections can be made vivid by considering Loar's notion of "singular visual directedness", which he illustrates with an example of successive lemon-experiences:

Suppose some indistinguishable lemons are one after the other brought to my visual attention. The lighting, the position of my eyes and so on, are all held constant. I am asked to think something about each lemon in turn, say, 'that's yellow'. Afterwards I am told that some of the apparent lemons were hallucinations.

... Those visual demonstrative concepts, and the perceptions that underlie them, are all *singularly visually directed*. (2003: 302)

The example might be simplified and diagrammed as follows:

<i>time</i>	t1	t2	t3
<i>content</i> ⁷	ℓ1	ℓ2	ℓ3
<i>object</i>	lemon 1	lemon 2	

The experience at t1 has content ℓ1 and lemon 1 is the object perceived; analogously for the experience at t2. The experience at t3 has content ℓ3, but no lemon is perceived: it is a hallucination of a lemon. In all three cases, the predicative element of the representational content is the same, namely *is yellow*. I will assume (I think in accordance with Loar’s intention) that the subject is not aware of the switch from one lemon to another, so there is no subjectively accessible difference of phenomenal character.⁸

Loar says that in all three experiences we have singular visual directedness, making plain that singular visual directedness does not require an object: it’s enough that all the experiences “purport to pick out some object visually” (302).⁹ Although experiences that differ in radical ways might share the property of singular directedness, it seems that for the lemon cases Loar has in mind something much more specific: the very same narrowest type of visual directedness is present in all three cases, in that the phenomenal character of all three both seems and is the same.

Loar clearly supposes that perceptual content will contain both singular and predicative elements. What should we say about the representational content of the envisaged singular element? I argue that there is no satisfying answer.

The two options worth discussing both treat it in a uniform way, as in each case the same or in each case different:

⁷ Strictly, the topic on this row is the representational *vehicle* at the relevant time. Content is a semantic property of a vehicle. Since in this case the candidates for content are lemons, it would be more accurate to regard the third row as specifying content. This way of putting it will sound strange to some ears, so at this point I adopt the more common idiom.

⁸ The formulation is intended to make room for the possibility that some experiences with subjectively indistinguishable phenomenal character in fact differ in their phenomenal character (think of phenomenal sorites).

⁹ Loar’s “all” is implicitly qualified by “concepts”, whereas in my presentation I make it sound as if it is qualified by “experiences”. This adjusts for my emphasis on the distinction between perceptual and conceptual systems of representation, detailed in §4 below. For language, the notion of purporting to refer can be explicated in a partially syntactic way: an expression purports to refer if it may properly occupy the kind of position in a sentence occupied by expressions that really refer. This explanation is not available if perceptual content lacks sentential structure. A functional role explanation is required.

a) the singular content is unchanged through t1 to t3.

b) the singular content is different at each of the three times.

Option (a) cannot be correct: evidently the three cases have different correctness conditions, and if this is to be determined by content, there must be variation in singular content, since the predicative content is unchanged. Suppose that the singular content throughout is that of e1, and so is, or refers to, lemon 1, and that this remains unchanged through the experiences. Then e3 would refer to lemon 1 and would not be a case of hallucination. Similar absurdities result whichever choice one makes of supposedly unchanged singular content.

Option (b) is more promising. It can do justice to the intuition that in the first two cases, the subject's perceptual system is working correctly: ℓ1 and ℓ2 are different, the first having lemon 1 as its referent, the second lemon 2, and both are correctly represented as yellow. And it can ensure that the experience at t3 is not correct, since ℓ3 has no referent. These results seem to me just what is wanted. But how could they be reflected in singular content? What mechanism explains the convenient but undetectable updating of the content of the supposed singular representational vehicle over the three cases?

External determination of content is, I shall assume, clearly possible in some cases. Oscar on earth refers to water by using "water", but Twin-Oscar on Twin Earth refers to twater by using a word that sounds just like the earth-word "water". Their phenomenology may be subjectively the same, even though their states involve different singular content. Is this a good model for the option (b) treatment of Loar's lemons?

The analogy is not exact. In Twin Earth cases the subjects are aware of a vehicle of content: a word, as I just presented the case, and there is a demand for determining what it means. The demand is met by providing a content that best explains subjects' relations to the world and their linguistic and cognitive activities. By contrast, in cases supposedly involving singular perceptual content there is no detectable vehicle of singular content, for such a vehicle would detectably change in the different cases; and the explanatory burden has already been discharged by the *perceiving* of the object, a causal-explanatory relation.

In the previous paragraph, I assumed that option (b) would register different contents by different vehicles. But there is a version of this option on which there is a single vehicle whose

content changes over time, perhaps in the way that demonstrative *words* can be regarded as having changing content, changes which in special cases may escape the conscious awareness of the subject. It is unclear to me what general theoretical structure would recommend this way of describing the lemons case. The alleged changing content plays no role in shaping either phenomenology or action.

On either version of option (b), the subject's perceptual state undetectably changes across the three times. No physical mechanism can explain the change, since perception-relevant energy patterns falling on the subject are unchanged, as are any perception-relevant activities in the subject. It is part of the way the example is set up that there are changes in the objects of the subject's perceptual states, and that these changing objects do not affect the subjectively accessible character of the states. Singular perceptual content would add an undetectable internal reflection of the external changes. But this is entirely superfluous from the explanatory point of view. It is given that the objects of perception change while the phenomenology is unchanged. The changes in the objects mean that, in a sense, the subject is in different states on the three occasions. But all this sameness and difference is straightforwardly accommodated by the causal facts, without alluding to singular perceptual content. Such content is theoretically unnecessary, and so does not exist.

In §4, I'll stress that I have no objection to singular *conceptual* representation, and I'll contrast this with supposed singular *perceptual* representation. In §5, I'll suggest a direct reason for being suspicious of singular perceptual content, one that is not premised on its theoretical dispensability.

4 Differences among representational systems

The broader picture I wish to paint is that whereas individual identity is readily available in the contents of thoughts, it is not available in the contents of perception. Hence there can be no singular perceptual content. In thought, we can readily make sense of the distinction between qualitative and numerical identity. This is not available in perception, whose

representational format is nonconceptual.¹⁰ The individuality or particularity of perception is fundamentally generated by causation and not by content.

Perceptual representational systems differ from conceptual ones in at least three ways¹¹: they have different expressive resources, different structural constraints and different possession conditions. As an example from the first category: conceptually, but not in visual perception, we can represent the content *is wholly invisible*. This shows that there is conceptual content that cannot be expressed by a visual perceptual content. The converse also holds: visual perceptual content makes very fine grained color discriminations. Perhaps the content of a perceptual experience is expressible by *red₂₂*, where “red₂₂” is a dummy expression standing in for some specific shade for which neither you nor I have a concept. This shows that there is perceptual content that is not expressed in the normal conceptual repertoire (though perhaps it is available to color experts). More generally, determinables (like *red*) rarely¹² feature in perceptual content unless some more determinate shade does, but there is no analogous connection for conceptual content; and many perceptually accessible determinates (like red₂₂) rarely or never feature in conceptual content. Perceptual content has an analog-based “richness” (amount of detail) that is hard or impossible to match in a more “digital” conceptual system. There are endless further differences: perhaps some logical words like “not” and abstract words like “justice” express conceptual content that cannot be literally expressed perceptually; perhaps contradictions, easy to express conceptually, cannot be represented perceptually¹³; and so on.

Differences of structural constraints were noted by Berkeley. You normally cannot visually perceive the color of an object without thereby having shape-related perceptual content¹⁴, but you can say or believe that something is a certain color, without saying or believing anything

¹⁰ Compare Pylyshyn (2007: 46): “true individuation requires concepts and conditions of identity”.

¹¹ These and similar differences are noted in the psychological literature, e.g.: “perception and conception differ in content, representational format, and methods of processing” (Mandler 2004: 45).

¹² The unusual cases are those in the visual periphery.

¹³ Unless examples like the waterfall illusion do really provide contradictory perceptual content, which is open to doubt.

¹⁴ In abnormal cases one might perceive the color of a Ganzfeld without perceiving an object or a shape; or one might be so close to a large object that it filled one’s visual field, thus depriving one of shape information but not of chromatic information. Berkeley connected color not to shape but to extension, and so his thesis does not require the restriction to normal cases.

about its shape. You may refer conceptually to a color using a semantically complex demonstrative phrase like “that color”, or “the color of those curtains”, but a visual experience of a color may have no matching complexity.¹⁵

Perceptual content has a high degree of involuntariness: open your eyes when there is light in your environment and, if you are sighted, you will have contentful visual experience whether you want to or not. Conceptual content, by contrast, involves some degree of voluntary manipulability. Although concepts are sometimes exercised involuntarily, as when one cannot help understanding words spoken in a familiar language, or when thoughts arise unbidden, having a concept requires you to be able at least sometimes to exercise it at will in the formation of thoughts. A possibly over-demanding version of this feature of concept possession is marked by Evans’s Generality Constraint: “if a subject can entertain those thoughts [that *a* is *F* and that *b* is *G*], then there is no conceptual barrier, at least, to his being able to entertain the thought that *a* is *G* or the thought that *b* is *F*” (1982: 102). Some degree of conceptual manipulability is an essential source of conceptual richness, making possible our most distinctive cognitive activities, like remembering, planning, premeditation and fiction.

The differences between perceptual and conceptual systems show that we cannot simply project features of conceptual systems onto perceptual ones. In particular, we cannot use the existence of conceptual singularity, which I do not dispute, to infer that there is perceptual singularity. In the next section, I’ll suggest that there could not be perceptual singularity: that is, there could not be perceptual representational vehicles with singular content.

5 Perceptual discrimination and content

There can be perceptually indistinguishable duplicates: two peas, two grains of sand, a genuine banknote and a brilliant forgery. To say they are perceptually indistinguishable is to say that, under the same conditions, a perceptual system, regarded as a causal mechanism, will respond to perceiving the one in just the same way as it would respond to perceiving the other.

¹⁵ “I can say that A and B are of different size without saying how much they differ in size or which is larger, but I cannot picture A and B as being of different size without picturing one of them as larger and indicating, roughly, how much larger it is. Similarly, if a yellow ball is situated between a red ball and a blue ball, I can state that this is so without revealing where (on the left or on the right) the blue ball is. But if this information is to be communicated pictorially...[e]ither the blue or the red ball must be pictured on the left.” (Dretske 1981: 137–8)

If they are placed side by side, they can be seen to be two, but perception cannot reveal which is which. When just one is perceived, everything of which the subject can be aware in perception is just the same as if the other had been perceived. The senses cannot discriminate the perceptually indistinguishable. This is not a fault in the system. Were it a fault, we could at least imagine how it could be fixed, but we cannot imagine a fix. The conclusion I draw is that to attribute singular content to a perceptual state would go beyond anything the state is capable of representing.¹⁶

A non-empty singular concept refers to just one thing. If it is applied to something other than that thing, the application is erroneous, however similar the object to which the concept is applied is to the object to which it refers. There is no analog for percepts. If there were singular percepts resembling singular concepts, then the perceptual system would be in error in responding to a duplicate in just the same way as to the genuine referent. But there would be no *perceptual* error in such a case, even if a false belief were formed. In the first two time intervals of the lemon examples, the qualitative identity of the lemons seen, together with the sameness of phenomenal character, makes it impossible to regard either case as one of incorrect perceptual functioning. True, the perceptual system failed to notice that lemon 1 had been replaced by lemon 2. But, according to the example, the switch was made undetectably, so there is no perceptual error. On the basis of the experiences, the subject may well form the false belief that the lemon she saw at t1 is numerically identical with the lemon she sees at t2.¹⁷ The belief is a conceptual structure, and employs representational resources not available to perception.

Perceptual content reflects how things appear, and indistinguishable duplicates cannot appear different (when viewed under the same circumstances).¹⁸ This informally reprises (T5),

¹⁶ Compare Churchland's description of a perceptual processing network: "The network itself draws no such distinction, between a reappearing individual on the one hand and sundry instantiations of an abstract universal on the other" (2013: 74).

¹⁷ The converse mismatch is familiar from the Müller-Lyer illusion: false perceptual content but true belief content.

¹⁸ Something like this view has been in the literature for some time. For example, McGinn (1982: 51) writes: "the content of experience is not to be specified by using any terms that refer to the object of experience on pain of denying that distinct objects can seem precisely the same". And Davies: "where there is no phenomenological difference for a subject, then there is no difference in content" (1992: 26). This thought has often been taken to favor an existentially general view of perceptual content, but that argumentative strategy depends on the false presupposition that singular and existentially general are the only options for perceptual content.

and (T1) shows how we can get individual objects into the story without appealing to singular perceptual content. (T1) implements singularity causally, rather than representationally.

The senses are essentially detectors of qualities. That is what they are tuned to do. Although they seem to yield individual-specific information, enabling us to recognize people and places, they achieve this thanks to the fortunate rareness of indistinguishable duplicates. In this, the senses are like any other “identifying mechanism” we possess or even can imagine: fingerprints, passwords, passports, Vehicle Identification Numbers, DNA profiles, iris scans, etc. These are complex qualities, and the qualities alone are causally involved in the initial output of any detector. The mechanisms can only represent the qualities of the causally relevant stimulus, and they will, evidently, not distinguish indistinguishable stimuli. By contrast, genuine singular content can be correct of only one of a collection of duplicates. A singular term like the name “Jim” does not apply to Tim, Jim’s identical twin, even if they share fingerprints and iris scans.

When we want to find out which individual object an iris scanner is scanning, we need to consider the scanned object, and not just the scanner itself, which would be in the same qualitative state even if it had scanned an indistinguishable object. Qualitatively identical states with different causes count as “different”. So a scan of Tim’s iris is in that respect different from a scan of Jim’s, even if their irises are qualitatively indistinguishable. But it is a familiar fact about perception, equally true of scanners, bank note detectors, fingerprints, signatures, and so on, that the signal itself does not individuate its cause: causes are fixed qualitatively, not singularly, in that a duplicate of something with certain causal powers shares those powers.

The underlying idea can be expressed in the following principle:

- (T7) A representational system includes only such representational vehicles as have contents that the system could in principle discriminate.

Since perceptual representational systems cannot even in principle discriminate qualitative duplicate individuals, they cannot include vehicles having such individuals as their content; that is, they cannot include singular contents. This contrasts with conceptual representational systems, which can in principle make a distinction in thought between an object and a duplicate of it. (Of course, this does not mean that the distinctions are correctly made in practice.) Enlarging the point, conceptual systems can allow for the possibility of fool’s gold and twin

water, perceptually indistinguishable duplicates of gold and water (according to the usual pretense). But this distinction is not available within a purely perceptual system.

Although this position seems to me straightforward and commonsensical, it may seem unable to do justice to various important features of perception: our ability to bind, to track, to “individuate”, to recognize. A single object can seem to have distinct properties, like being red and being round. How can the properties be *bound* to a single individual without the exercise of singular content? A predator *tracks* its prey across minutes or hours, locked on to the same target. We *attend* to individual objects, and *individuate* our friends by sight or voice. Argos *recognizes* Ulysses. We see things as *to the left or ahead or heading my way*. Surely all these perceptual activities are made possible by singular perceptual content? We need to represent that it’s the same individual that is red and round, the same individual that was there and is now here, that it’s Tim, not Jim, that it’s Ulysses; and egocentric properties seem to presuppose the singular identification of a self or other suitable origin.

In the next section, I’ll show how those activities find a natural place within the picture I propose; and in §7 I’ll enter a qualification.

6 Apparent problems for the feature-expression account

(a) *Binding*: One task set by “the binding problem” is to explain how different perceived features can be “bound together” as features of a single object. Psychology and neuroscience try to find explanations of this ability. A philosophical question that arises, and which is directly relevant to this paper, is: can a subject bind different perceived features as features of a single object without exploiting singular perceptual representational content? For example, suppose (to take a standard kind of philosophical example) a subject is confronted by a round red thing and a green square one. She registers the features round, red, square and green. But how does her representational system allocate the first two features to one object and the second two to another? What makes for the difference between seeing something red and round together with another thing that’s green and square, and seeing something green and round and another thing that’s red and square?

A possible answer is that the subject's perceptual system exploits singular perceptual representational content. A linguistic analog is that the correct binding occurs by exploiting two perceptual analogs of singular terms, one suitably conjoined with the features red and round, to the other with the features green and square. It can be doubted whether this is genuinely explanatory. In any case, the feature-expression account gives a different and more straightforward answer.¹⁹

If a perceived object triggers a feature representation, that feature is thereby "attributed" to the perceived object in the straightforward sense that the correctness of the representation depends on whether or not the represented feature is possessed by that object. This notion of attribution follows the relation of perceiving:

- (T8) If a representation of feature *F* is caused by perceiving the *F*ness of *x*, then *F* is attributed to *x*.

This is all we need in order to resolve the binding problem without appeal to singular representational content. A cluster of features is "bound to" an object *x* just if the perceptual situation is one in which all those features are "attributed" to *x* in the sense of (T8) (compare Neander 2017: 235). In the example, red and round are attributed to one thing, green and square to another, since the former directly caused the representation of red and the representation of round (and neither the representation of green nor the representation of square), while the latter directly caused the representation of green and the representation of square (but neither the representation of red nor the representation of round). Some object can satisfy *is red* and *is round* and another can satisfy *is green* and *is square* without any object satisfying both *is red* and *is square* or any object satisfying both *is green* and *is round*.

¹⁹ Compare Pylyshyn (2007: 47): "we may need to distinguish the process of feature-clustering and individuating from the process of picking out and referring to objects." Many observations in Pylyshyn's book seem sympathetic to the picture I offer, but a central ingredient of his theory, indices or FINGSTs, which are, apparently, singular perceptual representations, run counter to it.

(b) *Numerosity*: According to Pylyshyn (2007: 25–7), perceivers can

(i) perceive more than one object at a time, even if the different objects are qualitatively identical, and in a way that makes their distinctness apparent²⁰

(ii) accurately estimate the cardinality of small sets of things (maximum 5), if suitably arranged, using a faster and more automatic mechanism (“subitizing”) than ordinary counting.

These abilities seem closely related, and might suggest the deployment of singular perceptual representations. For example, perhaps the twoness of simultaneously perceived qualitative duplicates is perceptually represented by two singular representational vehicles with distinct referents, and these representations are the constant elements deployed through tracking. And perhaps subitizing reveals just how many of these singular representations a perceptual mechanism can produce at one time (and this is the moral Pylyshyn suggests in the discussion just referred to).

Explanations that appeal to singular representations are unnecessary. They may also be unsatisfying. For example, Pylyshyn suggests that one possible explanation for subitizing is that the subject may be “counting the number of indices deployed” (2007: 26), where indices are singular representations. But explaining the ability to count objects by the ability to count indices seems unhelpful, since the latter ability stands equally in need of explanation. My main point, however, is that although these experiments reveal psychological mechanisms of great interest, an explanation in terms of singular representations is not required.

Relational features can be represented in perception, for example *is taller than* or *is to the left of*. (The egocentric aspects of locational content are discussed shortly.) As in the examples, some relations obtain only between distinct things (they are antireflexive). Nothing more is required for a perceptual system to register distinctness (that is, the numerical feature of twoness) than to represent an antireflexive relation. Applying (T1), the perceptual content of a perceptual experience *x is taller than y* is correct iff there are perceived objects *x* and *y* of which

²⁰ As mentioned in the previous note, Pylyshyn’s indices (FINGSTs) may be singular contents, or vehicles of singular contents, posited to underlie this achievement: “a central function of indexes is to select and refer to (or bind arguments to) several visual elements at once” (2007: 28).

it is non-accidentally true. The relevant content ensures that it obtains of x and y only if $x \neq y$. The relevant relation could simply be distinctness itself.

The system could also represent cardinality features directly, like *being three*, a possibility Pylyshyn considers as an alternative to the use of indices (2007: 24). The skill of subitizing would then be described as skill with the application of such cardinality features.

(c) *Tracking*: A predator engaged in tracking by sight needs to try to keep the prey in sight. Or, if it loses sight momentarily, needs to try to regain sight. The lioness is tracking a specific antelope. She takes care to continue to perceive it, or to regain sight of it if continuity fails. Does not this involve a continual process of reidentification? And does not this in turn require a sequence of singular contents, perceptually connected by the equivalent in perceptual content of the identity relation? Should not tracking be modeled on the schema: x is at location L , y is at L^* , and although $L \neq L^*$, $x=y$? This might be understood as involving singularity in the perceptual representation of x and y .

A perceptual content could be, simply, *moves*. If the content is correct on an occasion, it is correct of a perceived object, one that moves, and so “remains the same” through change of place. This feature is integral to the feature *moves*, so movement as such need not introduce singular contents. (For motion, there are familiar questions, not to be addressed here, about the relation between the temporal extent of what is perceived and the temporal extent of the perceiving.) This gives the right result for illusions of movement. The seemingly moving scrolling messages seen on some public illuminated screens (e.g. in transit systems) are not really moving: it’s just that different lights come on successively. This is properly regarded as an illusion of movement, since the perceived object does not possess the perceptual content *moves*, despite the fact that perception insists on attributing this content even in the clear presence of the belief that it does not apply. As in Müller-Lyer cases, the inconsistency of perceptual content with belief content does not make the former disappear.

Tracking requires more sustained attention than the minimum needed for a perception to count as of movement. Must not the tracker exploit a singular content for what he is tracking? This worry should already have been assuaged by the account of binding. Tracking consists in

successive binding of successive features. In the favorable case, a single object is perceived across changes in place. This is a fact about perception, but is not reflected in singular perceptual content.

One of Pylyshyn's surprising results is that perceivers are good at tracking small numbers of similar objects through dissimilar paths of motion. A full account of tracking needs to apply to these special cases: MOTs (2007: 34ff). Although this is a skill calling for psychological explanation, it gives no special reason to suppose that tracking involves singular content.

(d) Recognition: Strictly speaking, recognizing an individual as such is not a perceptually available achievement. Instead, the most perception can offer is registration of the content *familiar* or some determinate of this (familiar parent, or whatever). As employed in perceptual systems, this is a purely qualitative notion, correctly attributed to objects that share qualities (to some contextually determined degree) with previously encountered objects.

Perception can at best target qualitative identity, whereas recognition, properly speaking, targets numerical identity. Since qualitatively identical but numerically distinct pairs are rare, qualitative identity is normally a perfectly adequate substitute for numerical identity, and the majority of the things we daily encounter are fungible in the sense that we would not care whether they were numerically the same (shirt, car, computer) or perfect duplicates. That is the basis of most real-life forms of "identification", by fingerprints, iris scans, passwords, passports, or VIN numbers. It is rare for distinct individuals to have indiscriminable fingerprints, irises or VIN numbers. So, normally, recognizing a qualitative pattern is as good as recognizing an individual as such. But not always: there is a small but non-zero probability of identical iris scans belonging to distinct individuals, and someone may fraudulently associate a vehicle with another vehicle's VIN number.

What perception does offer, as a kind of substitute for genuine recognition, is a feeling of familiarity, qualitatively based. Ulysses strikes Argos as familiar: he exhibits familiar perceptible qualities. But Argos would have been fooled by a duplicate. His perceptual system could not be blamed, had he been fooled, for no amelioration of it that we could even conceive would have prevented the error.

Familiarity and genuine identification come apart in certain cases. One is the Capgras delusion, whose victims claim that their closest friends, even spouses and pets, are mere qualitative duplicates, and not the real thing. Numerical identity is mistaken for merely qualitative identity. Another is *dejà vu*, in which what is in fact an unfamiliar scene induces a sense of familiarity, which may lead the subject falsely to believe in a numerical identity. In normal cases, when we recognize our spouse or a friend, though our perceptual basis is merely qualitative, our belief in numerical identity is well-founded.

The claim is not that perception may confuse qualitative with numerical identity. On the contrary, perceptual content cannot access numerical identity and so in a sense cannot confuse it with anything else. That's why the experience at t2 in the lemons case is correct: the perceived object is indeed yellow. The switch to a new lemon, distinct from the one seen at t1, was undetected, but perceptual experience does not claim otherwise, and hence cannot in itself be accused of anything incorrect on this basis. A false singular *belief* may arise, that lemon 2 is lemon 1, but this is made possible by additional representational resources at the conceptual level, which contains genuinely singular concepts.

(e) Perceiver-dependent properties A needed refinement relates to the representation of perceiver-dependent properties, like *being ahead or to the left or over there or heading my way*. Do these not make implicit reference to individual things: the perceiver herself, a direction, or a location she occupies? This is a hard question for all theorists²¹, since we regard a subject's capacities to self-refer or to refer to a specific location as relatively high level ones, but perceiver-relative contents seem to be required for the most primitive kind of perceiver, even magnetotactic bacteria²². The solution is that the perspectival coordinates, like the

²¹ The puzzle is well expressed by Peacocke: "An animal without language can see something as moving away from it, can hear something as coming towards it. These perceptions have a *de se* content, and it is very implausible that their occurrence in any way presupposes the possession of concepts ..." (2015: 177). My solution to the puzzle associates such experiences with subject-involving correctness conditions, but without requiring the subject to deploy a representational vehicle designating the subject. This is much less demanding than the capacity to think of oneself as oneself. Peacocke's attributions implicitly have sentential complements, e.g. "that something is moving away", and in my view these entail concept-possession on the part of the subject of the attribution.

²² These contain linear chains of magnetosomes which act like compasses, enabling bacteria containing these structures to orient relative to bodies of magnetic rock.

perceived objects, are not specifically represented in perceptual content, even though they shape its satisfaction conditions. They relate to perceptual content as gas tanks relate to gas gauge content.

Suppose a subject S is facing in direction D relative to S's position, P, and that under these circumstances, S veridically sees an apple off to the left. On one view, S's experience must contain singular contents for S, D, and P. But that seems too demanding, even setting aside the difficulties we have already encountered in getting objects, as opposed to their duplicates, into perceptual content. Self-identification is a notorious problem, and it's hard to see how there could be any direct way of identifying locations, as opposed to exhibited qualities: the most direct perceptual identification of a location is in the first instance an identification of the qualities present at it. They can be specified less directly, in terms of relations to other locations or directions, like the perceiver's position or angle of gaze, but these specifications in turn appear to involve something singular.

Egocentric parameters can feature in the satisfaction conditions for S's experience, without featuring in its content, just as the gas tank enters into the satisfaction conditions for a reading of the gas gauge. Suppose someone, S, sees an apple off to the left. The feature-expressive content is *an apple off to the left*. The satisfaction condition could be expressed: S's content is correct iff true of *an apple off to the left* of S, relative to her position, P, and direction of gaze, D, at the time of the experience. Only the italicized portion of this correctness condition features in the perceptual content. There is obviously a lot more detail to be filled in. For example, if you rotate your head to the right, ahead is not your direction of gaze (which is to the right), but is a projection of a line connecting your spine and your sternum (in that order). But as in the simpler example, these elements may appear in a correctness condition for a perceptual state without having proprietary representation in that state. This view enables us to accommodate the phenomenon for which "de se contents" were designed, without exaggerating what is involved in located perception. Located perception is available to simple organisms, as it should be. Once again, causation rather than content does the job.

If perception could reveal numerical as opposed to mere qualitative identity, it would involve singular content. But perception cannot do this, any more than any of our specially

engineered “identifying” mechanisms can. Claims of numerical identity are inevitably based on matters of qualitative identity.

7 A qualification

A basic iris scanner is connected to an electronically operated door latch. When someone approaches the door, their iris is scanned. The scanner checks the scan for a match with the scans in its database. If it finds a match, it opens the door. Otherwise, the door remains closed.

A more sophisticated iris scanner has an additional layer of data: the scans in its database are associated with names. This scanner can keep track of who is in the building at any time. It can maintain a display of everyone who’s there. It doesn’t merely open the door, it also keeps an individual tally.

Suppose Tim and Jim have qualitatively identical irises, ones that cannot be distinguished by any scanner. If Tim’s iris has been authorized, the basic scanner will open the door for Jim. (For simplicity, let’s suppose that Tim is not in the building.) As I was considering such a case earlier, this is a scanner that has not made a mistake: it has correctly identified an iris pattern. It would have made a mistake if it had contained singular content specific to Tim as opposed to Jim. The sophisticated scanner, by contrast, will respond to Jim not only by opening the door, but also by indicating that Tim is in the building even if he is not. The sophisticated scanner has made a mistake, by deploying an incorrect singular content. So some scanners do deploy singular contents.

Should we think of perceptual systems as like the basic scanner or like the sophisticated scanner?²³ A “basic” perceptual system would be purely qualitative, with no singular content. Its failure to discriminate duplicates is not a fault in the system, for we cannot even imagine how a system of this kind could be repaired in such a way as to enable it to make the discrimination. But a different kind of system, a “sophisticated” perceptual system, would couple the basic system with one that would output singular representations, and these outputs could be assessed for accuracy according to whether these representations are correct. Perhaps our actual perceptual systems are sophisticated, deploying singular content in

²³ Thanks to Martin Montminy for pressing this question.

recognizing one's spouse, one's friends, one's car. This process reliably delivers singular truth, even if the system could be fooled by a duplicate.

When thinking about individual objects, the question may seem unimportant. One might continue to say, with maximal conformity with what I have argued so far, that the "perceptual system", *strictly speaking*, is purely qualitative but may get coupled to a distinct system, perhaps a concept-deploying one, so that the resulting hybrid permits the deployment of content of both kinds. This is like saying that the sophisticated scanner is not strictly speaking just a scanner, but a scanner coupled with a distinct system which links the scanner's qualitative material to singular facts. Alternatively, one might say that although basic perceptual content is qualitative, sophisticated perceptual content may involve deployment of concepts. The source of the concepts is regarded as part of a unified system which also delivers the qualitative material. This is like saying that the sophisticated scanner is really a special kind of scanner, rather than a hybrid of a scanner and something else, but that scanners can do more than respond merely to qualities. Perhaps there is no important difference between these descriptions.

When one considers other things that we expect from perceptual systems, we find ourselves on a familiar philosophical battleground. I have said nothing about the features whose expression plays the central role in defining correctness for perceptual states in (T1). But many of these features are themselves subject to the argument from duplication. There could be something that is perceptually indistinguishable from a lemon but which is not a lemon; for example, it is a botanically unrelated but superficially indistinguishable fruit. If the earlier argument about singular content was good, this would entail that a content corresponding to *is a lemon* cannot be a perceptual content. Visual perceptual content would be confined to non-twin-earthable qualities, perhaps just the "sensible qualities", colors, shapes and perhaps some others. On this view, the perceptible qualities would be those for which there could be no perceptually indistinguishable duplicates. But this is a controversial claim. Many theorists take a much broader view of the range of nonsingular content that can be deployed in perception (for example, Schellenberg 2011). The broad view may seem to undermine what I have said about singular content. Given that perceptual systems can deploy contents for natural kinds, even if

they could not distinguish them from distinct natural kind duplicates (fools' kinds), why should not perceptual systems deploy singular contents?

The example of the basic and the sophisticated scanners may serve to moderate and somewhat refashion the debate. The analog of the broad view of the range of perceptual contents sees the sophisticated scanner as in some sense a single system with more than one kind of content. The analog of the narrow view, counting as perceptual content only cases that are insensitive to distinctions among qualitative duplicates, insists that only the basic scanner is a scanner, properly speaking, and that the sophisticated scanner is hybrid, a scanner bolted to a different kind of system.

Can this really be a matter of great philosophical debate? Everyone agrees that the perceptual systems of infants are much less discriminating than those of adults, that perceptual systems require the training of experience, and that this is a highly complex and long-running process during which the range of content output by the organism on the basis of its perceptual experiences is enormously enriched. Suppose an information processing system that begins with only non-singular content ends with singular content. Perhaps, for example, we want to say that Argos does genuinely recognize Odysseus, which entails that he must be in a state that deploys a singular content. On one view, we should say that this means he has coupled a nonperceptual system to his perceptual system. On another, it means his perceptual system has been trained to deliver this content, on a merely qualitative basis. I can imagine reasons for preferring one description to another coming from developmental psychology or neural anatomy; but not from philosophy.

Similarly, suppose a subject's overall information processing system begins with some qualitative input from a lemon and ends up with content true just of lemons, and not of lemon look-alikes. How much does it matter whether the output is allocated to the perceptual system as such, or rather to a hybrid system part of which is genuinely perceptual and part of which is not? On the former option, it has to be allowed that it is not a content that belongs to the basic or most primitive part of the perceptual system, for reasons given in this paper. There has to be a distinction of levels. But it is unclear to me how there could be philosophical reasons to prefer one as opposed to the other of the competing views: that the different levels demarcate a

distinction *within* the perceptual system, or that they mark a distinction *between* the perceptual system and a system of some other kind. I am therefore happy to regard the thesis of this paper as that there is no basic perceptual singular content. This is consistent with both the narrow view that there is no singular perceptual content and the broad view that there is singular perceptual content but none of it is basic.

References

- Burge, T. (1991). "Vision and intentional content". In E. Lepore & R. Van Gulik (Eds.), *John Searle and his Critics* (pp. 195–213). Cambridge, Mass: Basil Blackwell.
- Churchland, P. M. (2013). *Plato's Camera. How the Physical Brain Captures a Landscape of Abstract Universals*. Cambridge, Mass: MIT Press.
- Davies, M. (1992). "Perceptual content and local supervenience". *Proceedings of the Aristotelian Society*, 92, 21–45.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, Mass: MIT Press.
- (1995). *Naturalizing the Mind*. Cambridge, Mass.: MIT Press.
- Evans, G. (1982). *The Varieties of Reference*. Oxford: Clarendon Press.
- Grice, H. P. (1961). "The causal theory of perception". *Proceedings of the Aristotelian Society, Supplementary Volume*, 35, 121-152.
- Loar, B. (2003). "Phenomenal intentionality as the basis of mental content". In M. Hahn & B. Ramberg (Eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (pp. 229–257). Cambridge, Mass: MIT Press/Bradford Books. Reprinted in *Consciousness and Meaning. Selected Essays by Brian Loar*, Katalin Balog and Stephanie Beardman (eds), Oxford: Oxford University Press (2017: 291–317).
- Mandler, J. M. (2004). *The Foundations of Mind: Origins of Conceptual Thought*. Oxford: Oxford University Press.
- McDowell, J. (1991). "Intentionality de re". In E. Lepore & R. Van Gulik (Eds.), *John Searle and his Critics* (pp. 215–225). Cambridge, Mass: Basil Blackwell. Reprinted in his: *Meaning, Knowledge, and Reality*, Cambridge, Mass: Harvard University Press, 1998: 260–274.
- McGinn, C. (1982). *The Character of Mind*. Oxford: Oxford University Press.
- Neander, K. (2017). *A Mark of the Mental*. Cambridge, Mass: MIT Press.
- Pautz, A. (2010). "Why explain visual experience in terms of content?". In B. Nanay (Ed.), *Perceiving the World* (pp. 254–309). Oxford: Oxford University Press.
- Peacocke, C. (2015). "Perception and the first person". In M. Matthen (Ed.), *The Oxford Handbook of the Philosophy of Perception* (pp. 168–180). Oxford: Oxford University Press.
- Pylyshyn, Z. (2007). *Things and Places: How the Mind Connects with the World*: MIT Press.
- Schellenberg, S. (2011). "Perceptual content defended". *Noûs*, 45(4), 714–750.
- Soteriou, M. (2000). "The particularity of visual perception". *European Journal of Analytic Philosophy* 8: 173–189.
- Tye, M. (2018). "How to think about the representational content of visual experience". In C. Limbeck (Ed.), *The Philosophy of Perception, Proceedings of the Austrian Ludwig Wittgenstein Society: De Gruyter*