

Theories of perceptual content and cases of reliable spatial misperception

Andrew Rubner

Rutgers University, New Brunswick, New Jersey, USA

Correspondence

Andrew Rubner, Rutgers University, New Brunswick, New Jersey, USA.

Email: Andrew.rubner@gmail.com

Abstract

Perception is riddled with cases of reliable misperception. These are cases in which a perceptual state is tokened inaccurately any time it is tokened under normal conditions. On the face of it, this fact causes trouble for theories that provide an analysis of perceptual content in non-semantic, non-intentional, and non-phenomenal terms, such as those found in Millikan (1984), Fodor (1990), Neander (2017), and Schellenberg (2018). I show how such theories can be extended so that they cover such cases without giving up their core commitments.

1 | INTRODUCTION

A widely held view has it that perceptual states are representational states. That is, they are states that represent the world as being some way rather than some other way, and can be evaluated as accurate or inaccurate.¹ According to this view, misperceptions are inaccurate perceptual states—states which represent the world as being some way other than it actually is.

Given the rudiments of this view, it is at least conceivable that there could be a perceptual state that inaccurately represents the world any time it is tokened under normal conditions. For example, imagine a type of perceptual state which represents angles as being 33° but is reliably caused by 30° angles and is never, at least in normal conditions, caused by 33° angles. Any subject capable of tokening this perceptual state will suffer a misperception when they token it under

¹ For some criticisms of this view, see Brewer (2011), and Schellenberg (2014) for a comprehensive reply.

normal conditions. Let us call this a case of *reliable misperception*.² In what follows, I will show that there are actual cases of reliable misperception and that they cause trouble for theories of perceptual content which attempt to provide an analysis of what it is for a perceptual state to represent what it does (rather than something else) in non-semantic, non-intentional and non-phenomenal terms (for versions of such theories see, among others, Millikan (1984), Fodor (1990), Neander (2017), and Schellenberg (2018)).³ But more importantly, I will show how such theories can account for such cases without giving up their core commitments.

In section 2, I turn to a discussion of reliable misperceptions. I first provide a precise characterization of reliable misperception, and then, by appeal to some psychophysical experiments, argue that such cases are actual. The cases I focus on have a structure similar to the angle case discussed above. They show that there is a perceptual state R that represents a property F but that R tokens are not nomically correlated with F s but rather with G s, where G is in some sense approximately F (viz., in the sense that 30° angles are approximately 33° angles). In section 3, I leverage the cases discussed in section 2 to criticize some otherwise plausible and well-developed theories of perceptual content — namely, those which try to give an analysis of what it is for a perceptual state to represent what it does in non-semantic, non-intentional and non-phenomenal terms. In section 4, I develop a way of amending such theories of perceptual content so that they can handle the relevant cases of reliable misperception. To do this, I formulate two additional conditions that can be adopted by such theories. Roughly, the first states that if a perceptual state R represents a property F , then there is some property G such that G is nomically correlated with R and approximates F . An account of approximation in terms of a margin of error is provided. The second deals with why such approximate nomic correlations have arisen in the first place, rather than the perfect correlations our folk psychology may expect. Section five address some objections.

Before diving in, four preliminary remarks are in order. First, I will only be concerned with the perception of primary qualities. More specifically, I will focus my discussion on the perception of spatial magnitudes — properties like length, distance, and aspect-ratio (shape). As Peacocke (2015) reminds us, “[t]he perception of magnitudes is essentially involved in what is arguably the most basic kind of perceptual content, [...] scenerio content. The scenerio content of a perception is the way the perception represents the world in the immediate spatio-temporal environment of the subject as being” (p.378f.). Furthermore, my focus on magnitudes respects a distinction between magnitude classes, e.g., distance, and instances of magnitudes, e.g., a distance of 3 meters.

Second, I will take it for granted that perceptual states represent properties like *being 3 meters* and *being 30°*.⁴ In other words, I will take it that the vocabulary used to describe magnitudes is the same as the one used to describe the contents of perceptual states that represent magnitudes.

Third, when talking about the content of a particular perceptual state I will sometimes use *content-quotes*: $\langle \dots \rangle$. For example, a perceptual state R_F which represents an individual a as

² This informal characterization departs from others in the literature. For example, Mendelovici (2016) informally characterizes reliable misrepresentations, and hence reliable misperceptions, as representational states that get “things wrong in the same way all the time” (p.57). This is slightly stronger than her original characterization (also seemingly supported by Green and Rabin (2019, p.9)) that reliable misrepresentations are representational states which get things wrong in the same way all or most of the time (Mendelovici 2013). Artiga (2013) has pointed out that this stronger characterization is better suited for the cases she is concerned with. Likewise, my informal characterization is supposed to capture cases that I find worrisome for the theories of perceptual content to be discussed. A more precise characterization is provided in §2.1.

³ Other views put forward in the same spirit are found in Dretske (1981,1988), Papineau (1987), Price (2001), Shea (2018), and Schulte (2018). It is worth mentioning that some of these views are designed to handle not only perceptual states but other mental states as well.

⁴ If one likes one can rework such contents so they are unit free (see, e.g., Peacocke 1989, p.300ff.).

having a property F that is tokened at a time t will have the content <that a is F > at t (and so, I will be also assuming from here on that the contents of perceptual states have a predicational-demonstrative structure).^{5,6} Moreover, I will use content quotes to refer to perceptual state types. For example, “ R_F ” and “< that a is F >” denote the same type of perceptual state, namely, all and only those atomic perceptual states which represent an individual a as having the property F . And, I take it that perceptual states are type-individuated by their contents.

Last, I will assume that a perceptual state R_F that represents an individual a as having the property F is accurate if and only if a has F .

2 | RELIABLE MISPERCEPTION

Our first order of business is to discuss cases of reliable misperception and their significance. In this section I will provide an abstract characterization of reliable misperception (§2.1) and then provide some actual cases, gleaned from vision science, of reliable misperception (§2.2). The reader may find it helpful to move back and forth between these two subsections.

2.1 | Characterizing reliable misperceptions

To a first approximation, reliable misperceptions are perceptual states which represent the world inaccurately any time they are tokened under normal conditions. But we can make the notion of a reliable misperception more precise using the notion of a positive nomic correlation, where a state A is positively nomically correlated with a state B just in case the probability that A occurs *given that* B occurs is greater than the unconditional probability that A occurs, and that this is either because A states cause B states or because A states and B states have a common cause. More formally, where $p(\cdot)$ is a probability function defined in the normal way

- (PN) A state A is positively nomically correlated with a state B if and only if
- (i) $p(A|B) > p(A)$
 - (ii) (i) holds in virtue of the fact that either A states cause B states or that A states and B states have a common cause.⁷

An example will help clarify, fire-y states are positively nomically correlated with smoke-y states because the probability that there is fire present at time t is greater when there is smoke present at t than when there is no smoke present at t , and we take it that this is due to the fact that fire causes smoke. Last, note that we can easily define what it is for two states to be nomically correlated without qualification by replacing (PNi) with

- (i') $p(A|B) \neq p(A)$.

⁵ See Burge (2010) for critical discussion of this assumption.

⁶ I proceed this way to fix ideas. My arguments will go through even if the contents of perceptual states are conceived of differently.

⁷ (ii) captures the idea that the correlation is nomic. There is no need to read “cause” here as requiring that, if a state A causes a state B , then A necessitates B .

In our case we are concerned with positive nomic correlations between two kinds of states. One is the state of some F being in the range of a subject S 's sensory receptors. So, for example, in the case of visual perception, this would be the state of some F being in the line of sight or periphery of a subject S . The other kind of state is an R_F tokening —the tokening of a perceptual state that represents individuals as being F . So, to say that F s are positively nomically correlated with R_F tokens is to say that, for any subject S , $p(\text{some } F \text{ is in the range of } S\text{'s sensory receptors} \mid R_F \text{ is tokened by } S) > p(\text{some } F \text{ is in the range of } S\text{'s sensory receptors})$, and that this is because F s in the range of S 's sensory receptors cause S to token R_F or because whatever causes S to token R_F also causes F s to be in the range of S 's sensory receptors.

Before moving forward, it is important to note that when verifying whether a perceptual state R is positively nomically correlated with individuals which exemplify a property F we only care whether the relevant causal relation holds under certain background conditions. For simplicity, we can say that we only care if a particular stimulus (e.g., an individual exemplifying some property) is a normal cause of a perceptual state R . Let us say that a normal cause of a perceptual state R is a stimulus σ which would cause a tokening of R in a neurologically intact subject in their normal environment (natural habitat, ecological niche, etc.) and various enabling conditions hold (lighting conditions are of the right kind, etc.). In other words, if σ is a normal cause of R , then R would be caused by σ under normal conditions. Further, by saying that a stimulus σ is a normal cause of a perceptual state R , I am not saying that nothing other than σ normally causes R . Rather, I am saying that among the normal causes of R is σ . Nor do I think it is necessary that if σ is a normal cause of R , then the probability that a perceivable property F exemplified by σ is present when R is tokened is high. It is perfectly plausible that even if a stimulus σ that exemplifies a perceivable property F is a normal cause of a perceptual state R the $p(\text{some } F \text{ is in the range of a subject's sensory receptors} \mid R \text{ is tokened by the subject}) < .5$. This may occur if there are many stimuli which normally cause R .

We are now in a position to provide a precise characterization of reliable misperception.

(RM) A perceptual state R is a reliable misperception relative to a property F if and only if, for any subject capable of tokening R , R represents individuals as being F and it is not the case that F s are positively nomically correlated with tokens of R .⁸

Recall our first pass at characterizing reliable misperception. We said that reliable misperceptions are perceptual states which represent the world inaccurately any time they are tokened under normal circumstances. (RM) captures our informal gloss of reliable misperception nicely. To see why, suppose a perceptual state R satisfies (RM). So, there is no positive nomic correlation between R and the property F it represents individuals as being. From here we infer that F s are not a normal cause of R tokens (otherwise we would expect there to be a positive nomic correlation between them). Thus, given R will only be tokened accurately if it is caused by an F (in the right way), and F is not a normal cause of R , R will not be tokened accurately under normal circumstances.

⁸ Another formal characterization can be found in Menedelovici (2016, p.62; cf. p.61, fn.4, and p.70, fn.9). My characterization is more geared towards the perceptual case, whereas Mendelovici's is concerned with all kinds of mental states that misrepresent. The similarities and differences between the two characterizations are not of relevance here, but see §5 for where Mendelovici may push back at my characterization.




Figure 1 Example of a behavioral task used in Nundy et al. (2000).

2.2 | Actual cases of reliable misperception

Let us move on to discuss some cases of reliable misperception. I discuss two cases. They are chosen due to their simplicity and because they are well understood empirically. However, I take it that many other cases could have been focused on in what follows.⁹ (See McLaughlin (2016) and Green and Rabin (2019) for a discussion of how potentially widespread reliable misperceptions are.)

Case 1. The first case I want to focus on deals with angle perception. At a high level of generality this phenomenon amounts to the fact that “the subtense of any acute angle is seen as being somewhat larger than the measured angle of the stimulus, whereas the subtense of any obtuse angle is seen as being somewhat smaller” (Nundy et al. 2000, p.5592). Nundy et al. (2000; see also Carpenter and Blakemore (1973)) provide psychophysical evidence that neurologically intact subjects do indeed overestimate the size of acute angles, and underestimate the size of obtuse angles. Nundy et al. (2000) show this by having subjects participate in several different tests in which they “were asked to adjust the orientation of a line in relation to an angle that varied in subtense” (2000, p.5594). They used multiple tests in order to “minimize any confounding variables that might influence perception” (2000, p.5594). In one test used, for example, “subjects were asked to rotate a test line until it appeared collinear with the indicated arm of the inducing [i.e., stimulus] angle” (2000, p.5594, see Figure 1). If subjects were misperceiving the angle then they should place the test line either above the arm of the inducing angle or below it. The degree of misperception is measured by the disparity between where subjects actually placed the test line and where it would have been placed if the subject had successfully made the test line collinear with the relevant arm

⁹Mendelovici (2013, p.424 and 2016, p.70, fn.9) has suggested that the cases to be discussed should not count as cases of reliable misperception. I address this line of thought in §5.

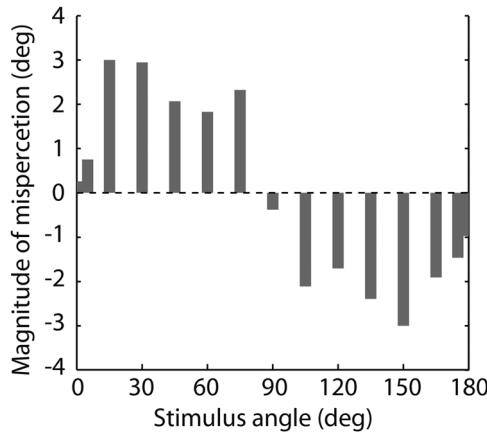


Figure 2 Summary of Nundy et al.’s (2000) data from Howe and Purves (2005).




Figure 3 Data from Hibbard et al. (2012). The data can be described by the following equation, where R is slant, and g is the gain at which subjects use info from depth cues, k is a constant, and \hat{A} is perceived aspect ratio: $\hat{A} = k(\cos(gS)) = .87(\cos(.66S))$.

of the inducing angle. It is important to note that subjects were not asked to assign numerical values to angles in any of the tests given.¹⁰

The findings of Nundy et al.’s experiment showed that subjects judged acute angles to be slightly larger than they actually were and obtuse angles to be slightly smaller than they actually were. For example, as can be seen clearly in Figures 2 and 3, subjects judged 30° angles to be about 33°, and angles of about 150° to be about 147°. Thus, Nundy et al.’s experimental evidence shows

¹⁰ One may try to argue that subjects will not misperceive angles in this way when viewing them in more realistic conditions (whatever those may be), as opposed to the artificial conditions in which Nundy et al.’s (2000) experiments take place. However, the kind of misperception that Nundy et al. (2000) are investigating underlie many classical persistent illusions, such as the Zollner and tilt illusions (Howe and Purves 2005, p.38f.). And, it is plausible that such illusions will persist in more realistic settings, as many persistent illusions of a similar nature do. See Howe and Purves (2005, p.5f.) for discussion.

that perceptual states that represent angles as being 33° are positively nomically correlated with angles which are 30° and not positively nomically correlated with angles which are 33° . We draw this inference on the basis that a natural explanation of why Nundy et al.'s subjects behave the way they do is because angles of 30° cause perceptual states which represent angles of 33° under normal conditions. This in turn gives us reason to believe that, where a is an angle,

- (1) For any subject S , $p(\text{an angle of } 30^\circ \text{ is in } S\text{'s line of sight or periphery} \mid \langle \text{that } a \text{ is } 33^\circ \rangle \text{ is tokened by } S) > p(\text{an angle of } 30^\circ \text{ is in } S\text{'s line of sight or periphery})$.

And, likewise, that a natural explanation of why subjects behave as they do when they are shown 33° angles is because angles of 33° cause perceptual states, when the subject is operating under normal conditions, which represent angles as being 36° . This in turn gives us reason to believe that

- (2) For any subject S , $p(\text{an angle of } 33^\circ \text{ is in } S\text{'s line of sight or periphery} \mid \langle \text{that } a \text{ is } 36^\circ \rangle \text{ is tokened by } S) > p(\text{an angle of } 33^\circ \text{ is in } S\text{'s line of sight or periphery})$.

Indeed, given that (2) is true and the results of Nundy et al.'s experiment, we have no reason to believe that

- (3) For any subject S , $p(\text{an angle of } 33^\circ \text{ is in } S\text{'s line of sight or periphery} \mid \langle \text{that } a \text{ is } 33^\circ \rangle \text{ is tokened by } S) > p(\text{an angle of } 33^\circ \text{ is in } S\text{'s line of sight or periphery})$.

Indeed, we have good reason to believe that (3) is false, as it is difficult to see how it could be compatible with our reasons for believing the truth of (2). And so, we have a case of reliable misperception. Similar reasoning shows that other cases of angle (mis)perception uncovered by Nundy et al.'s experiment are exemplars of reliable misperception.

Case 2. The second case we will look at deals with shape perception. More specifically, this case deals with the fact that circles are misperceived as slightly elongated ellipses, and that slightly shorter than wide ellipses are perceived as being circles. More specifically, Hibbard et al. (2012) showed that subjects judge, on average, ellipses with an aspect ratio of .87 to have an aspect ratio of 1. Moreover, Hibbard et al. (2012) claim that the misperception is a misperception of aspect ratio that arises due to the more general fact that perceived line length is a function of orientation insofar as two lines of the same physical length may be perceived as differing in length if oriented differently (Figure 4). For example, that illusion tells us that two lines of the same length will be perceived as differing in length given one is oriented vertically (the line that is perceived as longer) and the other horizontally. Because of this, it is reasonable to suppose, given Hibbard et al.'s experimental results, that circles will not be perceived as being circular, but rather as slightly elongated ellipses.

Hibbard et al.'s experiment used neurologically intact subjects viewing in *enhanced* cue conditions. That is, even when provided information about the slant of the surface via binocular disparity goggles, subjects still judged ellipses with an aspect ratio of .87 to have an aspect ratio of 1. This is significant because, trigonometrically, one can infer the aspect ratio of a 3D surface projected on to a 2D plane if they are provided with information about the slant of the 3D surface and the aspect ratio of the projected image (see Figure 5). And, it is often thought that the visual system can make such an inference given the requisite information. In sum, Hibbard et al. report that they "we were able to find no binocular-disparity defined slant at which the effective slant




Figure 4 Perceived line length is a function of orientation. The vertical line is perceived as being about 13% longer than the horizontal line despite being the same physical length. The line about 30° from the horizontal plane is about 15% longer than the horizontal line, making it the line which makes the greatest departure from its horizontal counterpart. See Craven (1993) and Zhu and Ma (2020) for psychophysical evidence.



Figure 5 If x is the aspect ratio of the image projected on to a 2D plane from a 3D surface, then one can infer the aspect ratio of the 3D surface (A) given the slant of the 3D surface (R). That is, $A = x(\cos(S))$. Image from Geisler and Kersten (2002).

away from the frontoparallel plane was zero, such that an ellipse with an aspect ratio of 1 in the image appeared circular” (Hibbard et al., 2012, p.41). In other words, at no slant is an ellipse with an aspect ratio of 1 perceived as being circular.

Now, given the behavior of the subjects in Hibbard et al.’s (2012) experiment, it is plausible that circles do not normally cause perceptual states which represent individuals as being circular (i.e., as being elliptical with an aspect ratio of 1). Indeed we have reason to believe that circles normally cause visual percepts which represent individuals as being elliptical with an aspect ratio of 1.13, and that perceptual states which represents individuals as being elliptical with an aspect ratio of 1 are normally caused by elliptical individuals which have an aspect ratio of .87. This in turn gives us reason to believe the following two propositions, where a is an elliptical surface,

- (4) For any subject S , $p(\text{an ellipse with an aspect ratio of } 1 \text{ is in } S\text{'s line of sight or periphery} \mid < \text{that } a \text{ is } 1.13 > \text{ is tokened by } S) > p(\text{an ellipse with an aspect ratio of } 1 \text{ is in } S\text{'s line of sight or periphery})$.
- (5) For any subject S , $p(\text{an ellipse with an aspect ratio of } .87 \text{ is in } S\text{'s line of sight or periphery} \mid < \text{that } a \text{ is } 1 > \text{ is tokened by } S) > p(\text{an ellipse with an aspect ratio of } .87 \text{ is in } S\text{'s line of sight or periphery})$.

But no reason to believe that

- (6) For any subject S , p (an ellipse with an aspect ratio of 1 is in S 's line of sight or periphery | \langle that a is 1 \rangle is tokened by S) $>$ p (an ellipse with an aspect ratio of 1 is in S 's line of sight or periphery).

Thus, Hibbard et al. (2012) have provided us with another case of reliable misperception.

3 | WHY NOT BUSINESS AS USUAL?

It is time we spell out why these cases are problematic from the perspective of the current theories of perceptual content. Given the wide range of views on the market, it will not be possible to survey them all here (at least not without making this paper inordinately long). The goal in this section is more modest, it is simply to discuss three plausible and well-developed theories — viz., Fodor's asymmetric dependency theory, an informational teleosemantics, and Schellenberg's (2018) capacatism —and show that they have trouble handling the cases that we have been discussing.¹¹ Recall that this paper is supposed to be constructive, so we will also use the views discussed here to develop a positive proposal that can, in principle, be adopted by any view.

A remark of clarification is in order before I begin, however. It should be noted that the views to be discussed are compatible with the existence of at least some cases of normal misperception —cases of misperception that occur even when a subject's perceptual systems are not malfunctioning and viewing circumstances are normal. However, the cases which such views can handle have the following characteristic which makes them different from reliable misperceptions (which are a subset of normal misperceptions): they involve perceptual states which get things right a non-negligible amount of the time.

To demonstrate what I have in mind, take into account the case in which a creature has a perceptual state R_p that represents individuals as having the property *predator*. It may be beneficial, from an evolutionary perspective, for this state to be tokened by anything which *may* be a predator, since it is very important that it is tokened in the presence of a predator. But, what the last clause of the previous sentence tells us is that R_p is sometimes tokened accurately under normal conditions. Indeed, if the creature is to be evolutionarily successful it will need to be tokened correctly whenever predators are actually in the creature's line of sight or periphery (cf. Godfrey-Smith 1991).¹²

It is worth highlighting that Godfrey-Smith's classic (1991) description of this case would not constitute a case of reliable misperception by the lights of (RM). He describes it as follows: there is a subject S that has a perceptual state R_p that represents a property *predator* where the prior probability of their being a predator present in S 's environment is .2, so p (there is a predator here now) = .2. But, given certain background assumptions about how valuable it is to detect predators and how cheap it is to token R_p falsely, the conditional probability that there is prey here now given R_p is tokened is only slightly greater than the unconditional probability that there is prey here now, i.e., the p (there is a predator here now | R_p is tokened by S) = .35 (1991, p.719f.; how

¹¹ A consumer-based teleosemantics is addressed in §5. It has been argued that such views are able to handle cases of persistent illusions (such as the Ebbinghaus illusion) which bear striking similarity to the illusions discussed above. This potential objection is addressed in §5.

¹² Other cases of non-reliable normal misperception involving categorical perception can be found in Matthen (1988).

he comes upon this quantity is not at issue here). Now, by (RM), R_p is a reliable misperception if and only if, for any subject capable of tokening R_p , R_p represents individuals as being predators and it is not the case that predators are positively nomically correlated with tokens of R_p . The first conjunct on the left hand side of the bi-conditional holds by stipulation. So, if R_p is a reliable misperception, then, predators are not nomically correlated with R_p tokens and so either it is not the case that

- (7) For any subject S , $p(\text{there is a predator here now} | R_p \text{ is tokened by } S) > p(\text{there is a predator here now})$

or (7) is true but not in virtue of the fact that predators cause R_p tokens in subjects capable of tokening R_p or because whatever causes R_p tokens in subjects capable of tokening that state also causes predators to be in the presence of such subjects. By stipulation (7) is true. So, if this case is to count as a reliable misperception it must be because (7) does not hold in virtue of the causal facts described above. However, it is plausible that Godfrey-Smith (1991) is implicitly assuming that predators do cause R_p tokens (see his discussion around p.713). Indeed, it is unclear whether the case can be made plausible while denying this. And thus, I take it that any plausible re-description of the case will not count as a case of reliable misperception according to (RM).

Now, it will help to impose a uniform structure on the theories we will examine. Namely, each theory can be formulated as providing an explanatory sufficient condition for what it takes for a perceptual state R to represent an individual as being F . I will set aside what is meant by an “explanatory sufficient condition” and just assume each theory makes good on their promise to provide one. Moreover, the theories I will look at are restricted in three ways. First, they are non-circular insofar as the explanatory sufficient condition (or conditions) do not make appeal to representational or otherwise intentional (or semantic, etc.) terms.¹³ Second, such theories do not appeal to notions related to phenomenal consciousness.¹⁴ And third, such theories are atomistic insofar as they primarily assign content to atomic perceptual states and then derivatively assign content to complex representations “recursively by the sorts of techniques that are familiar from the construction of truth definitions” (Fodor 1990, p.58).¹⁵

In sum, then, the theories to be discussed will each take the following form:

- (#) An (atomic) perceptual state R represents an individual as being F if C , where C contains no phenomenal or representational terms.

It should be noted that not every theory discussed below is originally formulated this way. For example, some theories are formulated using a bi-conditional and other theories are not

¹³ I will not provide a defense of this constraint here, but a traditional motivation for answering this question in such a way comes from a reductivist perspective. For example, Fodor writes that “*aboutness* [and hence *representation*] surely won’t” appear on the “complete catalogue [...] of the ultimate and irreducible properties of things” (Fodor 1987, p.97). However, one may want to motivate such a methodology more modestly by pointing out that “any view of representational content must explain in virtue of what a mental state has a specific representational content” and that it would be “explanatory unsatisfying to” answer such a question by appeal to representational notions (Schellenberg 2020, p.757). Not everyone would agree, of course. See, for example, Burge (2010, ch.8). There is a sense in which Schellenberg (2020) is responding to Burge in the above passages.

¹⁴ This distinguishes such views from phenomenal intentional theories of perceptual content. See Mendelovici and Bourget (2014) for helpful discussion of how these views relate to the views at issue here.

¹⁵ The consumer-based account discussed in §5 drops this assumption.

formulated as providing an explanatory sufficient condition at all (e.g., Schellenberg's (2018) theory is not). That is okay for our purposes, since our goal is not to follow these theories to a tee, but rather to get a feel for them and their discontents in light of reliable misperception.

The first theory I will look at is an approximation of Fodor's asymmetric dependency theory. Here is Fodor's statement of his view (it is worth noting that Fodor's theory is developed with respect to all representational states which explains his example):

So, what the story about asymmetric dependence comes down to is that "cow" means cow if (i) there is a nomic relation between the property of being a cow and the property of being a cause of "cow" tokens, and (ii) if there are nomic relations between other properties and the property of being a cause of "cow" tokens, then the latter nomic relations depend asymmetrically upon the former. (1990, p.93)

Given the above quote, we can state Fodor's view using our notation as follows:

- (F) A perceptual state R represents individuals as being F if
- (i) there is a nomic relation between the property of being an F and the property of being a cause of R tokens; and
 - (ii) if there are nomic relations between other properties and the property of being a cause of R tokens, then such nomic relations depend asymmetrically upon the nomic relation cited in (i).¹⁶

To be clear, for Fodor, if there is a nomic relation that holds between the property of being an F and the property of being a cause of R tokens, then F 's cause R tokens is a counterfactual supporting generalization and consequently that there is a "covering" law that relates the property of being F to being a cause of R tokens (1990, p.93; Fodor puts "covering" in scare quotes). Although it is not clear what exactly Fodor has in mind here, I am assuming he is saying that if condition (Fi) holds for some perceptual state R and property F , then it is a law that

- (L) F s cause R tokens

and that, because of this, something like the following counterfactual is true:

- (CP) *Ceteris paribus*, if an F were to stimulate the sensory receptors of a subject S in the right way, then R would be tokened.

To show that Fodor's view (and the others) cannot handle cases of reliable misperception, let us focus on a particular case of reliable misperception. Namely, let us focus on R_{33° , the perceptual state that represents individuals as being 33° angles. Now, suppose that the antecedent of (F) is true for R_{33° . Thus, according to (F), we have:

- (8) There is a nomic relation between the property of *being a 33° angle* and being a cause of R_{33° tokens

¹⁶ Fodor also states the view with an extra clause, namely, that some R tokens are actually caused by F s (1990, p.121). Nothing hangs on this condition being elided. If anything, it raises more issues for Fodor.

which in turn implies

(9) 33° angles cause R_{33° tokens

which in turn implies

(10) *Ceteris paribus*, if a 33° angle were to stimulate the sensory receptors of a subject S in the right way, then R_{33° would be tokened.

However, since Nundy et al.'s (2000) experiments give us reason to deny that 33° angles are a normal cause of R_{33° , it is unclear why we should think (10) is true. To put it more succinctly, if (10) were true, we would expect there to be some conditions under which the counterfactual scoped under the *ceteris paribus* clause in (10) is true, but the best candidate conditions are those under which Nundy et al.'s experiment take place. And, as evidenced by Nundy et al.'s experiment, even in those conditions that counterfactual is false. So, we have no reason to believe that (10) is true. I take it that other cases of reliable misperception will end up posing similar difficulties. And thus, I take it that (F) cannot account for cases of reliable misperception.

Let us move forward. An informational teleosemantics may be formulated as follows:

(IT) A perceptual state R represents individuals as being F if R has the function to carry information about F s.¹⁷

Two parts of this theory require comment. First, the relevant notion of function for such theories is that of a *normal-proper* function. Such functions are meant to capture what an entity (e.g., the heart, kidney, a perceptual state) is *supposed* to do. For example, it is the normal-proper function of the heart to pump blood, because hearts are supposed to pump blood. Importantly, given that a normal-proper function is what an item is supposed to do, a token item that has a normal-proper function to ϕ can have that function without being able to ϕ . Further, the relevant notion of carrying information can be defined in terms of nomic correlations. Namely, we can say that a state B carries information about a state A just in case A states are nomically correlated with B states, i.e., just in case $p(A|B) \neq p(A)$ because A states cause B states or B and A states have a common cause (cf. Skyrms 2010, p.35f.; Shea 2018, p.75ff.; Ganson 2021, p.686f.). Moreover, let us say that

(I) A perceptual state R carries information about F s if and only if, for any subject S , $P(\text{some } F \text{ is in the range of } S\text{'s sensory receptors} \mid R \text{ is tokened by } S) \neq P(\text{some } F \text{ is in the range of } S\text{'s sensory receptors})$.

Again, suppose that the antecedent of (IT) is true for R_{33° . Thus, we have the following:

(11) R_{33° has the function to carry information about 33° angles.

The first observation I want to make is that, if (11) is true, then I think it must be true without it being the case that R_{33° carries information about 33° angles, as it is unclear how this could be

¹⁷Versions of this view have been defended by Neander (1995, 2017) and Schulte (2018), among others. See also Rubner (2022, p.5f.)

true given that Nundy et al.'s experiments provide reason for thinking that (3) (reprinted below) is false.

- (3) For any subject S , $p(\text{an angle of } 33^\circ \text{ is in } S\text{'s line of sight or periphery} \mid < \text{that } a \text{ is } 33^\circ > \text{ is tokened by } S) > p(\text{an angle of } 33^\circ \text{ is in } S\text{'s line of sight or periphery})$.

However, it is not impossible, at least according to many views of normal-proper function, for a type of item i to have a function to ϕ without any i tokens being able to ϕ . The challenge, rather, is to make this claim plausible. One immediate response, for example, would be to argue that (i) R_{33° tokens were once nomically correlated with 33° angles but that (ii) this changed due to a radical shift in the environment. The first conjunct would thus explain why R_{33° has the function to carry information about 33° angles and the second would explain why it no longer does. This response, however, does not work for either of the cases here. This is because a leading explanation for our misperception of angle size and aspect ratio deals with the fact that our visual system takes into account the statistical regularities found in non-carpentered environments that have been internalized over the course of evolution when inferring the 3D shape of objects from a 2D retinal image.¹⁸ So, claiming that these cases of reliable misperception arise because there was a radical shift in the environment seems to go directly against this kind of explanation for the etiology of these cases of misperception. Such explanations support the idea that we, as a species, currently have such misperceptions because we have been in the same environment long enough to internalize the regularities of that environment, which goes against the idea that we have such misperceptions because there has been a radical shift in our environment. That said, however, this kind of response may work for other cases of reliable misperception, but importantly, this would have to be shown once the details of the relevant cases are on the table.

Let us move on finally to Schellenberg's (2018) view. According to Schellenberg, perceptual content is constituted by a subject's employment of perceptual capacities that have the function to discriminate and single out particulars of a certain kind. Like the teleosemantic view presented above, I take it that Schellenberg is also making use of the notion of a normal-proper function insofar as token capacities can fail to fulfill their functions, albeit her preferred notion of normal-proper function is neither etiological nor relativistic (2018, p.35f.). So, whether a capacity has the function to ϕ is neutral on whether it succeeds in ϕ -ing on a particular occasion. The function of a capacity, for Schellenberg, is thus plausibly tracking what a capacity is supposed to do, not what it actually does on some given occasion.

More importantly, on Schellenberg's view employing a perceptual capacity constitutes a perceptual state that represents the kind of particular the capacity has the function of discriminating and singling out. For our purposes, we can think of kinds of particulars as properties. Moreover, Schellenberg's view does not take perceptual states to have an attributional structure (see her 2018, p.67ff.), but to maintain continuity I will treat as if they do. So, imposing our structure on to Schellenberg's view we have:

- (S) A perceptual state R represents individuals as being F relative to a perceptual capacity C_F and subject S if
- (i) S 's employment of C_F puts S in R ; and
 - (ii) C_F has the function to single out and discriminate F s.

¹⁸ See Howe and Purves (2005, ch.4) for a detailed discussion and support for this kind of explanation.

A similar problem arises for Schellenberg's view that arose for the teleosemantic theory above. Suppose that the antecedent of (S) is true for R_{33° . From this it follows that there is a perceptual capacity C_{33° that has the function to single out and discriminate 33° angles. But, given Nundy et al.'s experiment, it is unclear why we should think this function is ever fulfilled. If Schellenberg's view is correct, then in Nundy et al.'s experiments, subjects were employing C_{33° without 33° angles being singled out and discriminated by them, even when they were operating under normal conditions. Thus, it is unclear what conditions would need to hold in order for C_{33° to fulfill its function. But, this calls into question whether the capacity genuinely has the function it is said to have. So, the challenge is the same as it was for (IT).

Another issue arises that is more internal to Schellenberg's theory, which is also worth mentioning. Schellenberg provides the following possession condition for an arbitrary perceptual capacity:

A subject S possess a perceptual capacity C_α if and only if the following counterfactual is true of S : S would be in a position to discriminate and single out a particular α_1 , where α_1 is any particular of the type that C_α functions to discriminate and single out, if R were perceptually related to α_1 , assuming (i) S is perceptually capable (awake, alert etc.), (ii) assuming no finking masking, or other exotic cases obtains, and (iii) where R being perceptually related to α_1 means that (a) the situational features are such that α_1 is perceivable by S (good lighting conditions, etc.), (b) R has the relevant sensory apparatus that allows her to gain information about α_1 , and (c) S is spatially and temporally related to α_1 such that S is in a position to gain information about α_1 via her sensory apparatus. (2018, p.31f.)

Now, if we have a case of reliable misperception, then it is unclear whether a subject does possess the relevant perceptual capacity according to Schellenberg. To see why, it may be helpful to look at another case from §2.2. Since human subjects reliably misperceives 30° angles (or any acute angle for that matter), it may be argued that humans do not have a visual system (the relevant kind of sensory apparatus) that allows them to gain information about 30° angles and thus condition (iii,b) cited in Schellenberg's possession condition would go unsatisfied. But if subjects do not have the relevant perceptual capacity, then we do not have a story of in virtue of what the perceptual state R_{30° represents what it does on Schellenberg's view.

What I have taken myself to have shown thus far is that the three views just discussed face challenges when trying to account for the cases of reliable misperception we have been concerned with. I am not claiming that these challenges cannot be overcome with the current resources of those theories, but rather that there is no obvious way to overcome them with those resources. More generally, I think the issue is that these views seem to work best, or are most plausible, *only if* there is a nomic correlation between the relevant perceptual states and the properties that those states represent individuals as having, and I suspect this will also be true for other views that were not glossed here. Cases of reliable misperception are cases in which no such nomic correlation holds. So, instead of looking at more views which I take to be implausible given there are cases where no such nomic correlations hold, I will develop a way of amending already existing views so that they can handle cases, within certain limitations, in which a nomic correlation does not hold between the perceptual state and the property it represents individuals as having.

4 | A PROPOSAL

One can amend the views above by incorporating two additional explanatory sufficient conditions into them. Let us look at a sketch of the two conditions as distinct from any theory in particular, and then see how they can be incorporated into those theories.

The first condition states the relationship between the relevant perceptual state and the property it is nomically correlated with. More precisely, it states that the value of the magnitude being represented by the perceptual state approximates the value of the magnitude that the perceptual state is nomically correlated with.

(C1) *R* represents an individual as being *F* relative to a property *G* if *R* is nomically correlated with individuals that exemplify *G*, where both *F* and *G* are of the same magnitude-type \mathcal{M} . And, the determinate value of *G* approximates the determinate value of *F*; and (...).

(NB: The ellipse is to signal that more conditions will be needed to get an adequate theory of perceptual content.)

Approximation can be further cashed out in terms of a margin for error. More specifically, we can say that:

(A) A value v_1 approximates a value v_2 relative to a margin of error $m \geq 0$ if and only if $|v_1 - v_2| \leq m$.¹⁹

How is a margin of error determined? What the relevant margin of error is will depend on the case at hand. In general, however, we can get a feel for what the margin of error should be by looking at the reported degree of misperception by the empirical studies which confirm that our perception of magnitudes are reliably mistaken. For example, recall Nundy et al.'s studies of angle misperception. According to those studies, the most extreme misperception of angle was by about 3°. So, for cases of angle perception, it is plausible that the margin of error is 3°. This, of course, will not always be the case. In the examples of misperception of aspect ratio, we may set the margin of error of around .15, because that illusion arises from the more general phenomenon that (mis)perceived line length is a function of orientation and the largest illusion of that kind is of about .15 (Figure 4). Facts about margins of error, then, become empirical facts.

The second condition fills in a lacuna that the first condition cannot fill. It seeks to explain why a perceptual state reliably misperceives in the way it does (when the perceptual state does reliably misperceive). For example, it answers questions like: why do 30° angle stimuli give rise to percepts of 33° angles under normal conditions? Or, why are ellipses with an aspect ratio of .87 perceived as circles? To make the present proposal as general as possible, I will not claim that there is a single explanation for why a perceptual state reliably misperceives in the way it does. Rather, I will formulate the second condition using an open parameter *E*.

(C2) A perceptual state *R* represents an individual as being *F* relative to a property *G* if *E*, where *E* explains why *G*s are nomically correlated with perceptual states that represent *F*; and (...).

Filling in *E* would provide an explanation for why a perceptual state that is a reliable misperception goes wrong in precisely the way it does. Of course, for any given case *C*, *E* will have to

¹⁹ Worries that this notion of approximation is inadequate are addressed in §5.

be cashed out in non-semantic and non-intentional terms. We will see how this may go when looking at how a theory of perceptual content can be modified using our two conditions. Before moving to that, however, it is important to note that if we conceive of (C2) as being a condition that allows us to explain why a perceptual state reliably misperceives in the way it does, then it will be trivially satisfied in the case where R is not a reliable misperception. More precisely, we can say that

(%) For any E, R , if R is not a reliable misperception, then $E = \text{“}R \text{ is nomically correlated with the property that } R \text{ represents individuals as having”}$.

Before moving forward, it is worth making explicit that by introducing (C1) and (C2) into a theory of perceptual content one will be committed to fixing the contents of a perceptual state R_F relative to another property G (where it is left open that $G = F$). However, “ R_F ” (and so “ $\langle \text{That } a \text{ is } F \rangle$ ”) is still to be understood as denoting all and only those atomic perceptual states which represent an individual a as having the property F .

Now, without further ado, let us see how our conditions can be used to amend a pre-existing theory. To fix ideas, let us focus on how the informational teleosemantics discussed above can be amended using our conditions, and then see how such a theory is able to handle our problematic cases (how other theories may be amended is left to an appendix).

(IT') A perceptual state R represents an individual as being F relative to a property G if

- (i) it is the function of R to carry information about individuals which approximate F s;
- (ii) R is nomically correlated with individuals that exemplify G , where F and G are of the same magnitude-type \mathcal{M} , and the determinate value of G approximates the determinate value of F ; and
- (iii) E , where E explains why G s are nomically correlated with perceptual states that represent F .

Note that in the case where $F = G$ we can recover the core of (IT). That is, if $F = G$, then we do not have a case of reliable misperception. Thus (ii) will be trivially satisfied and the E appearing in (iii) will be filled out in accordance with (%). Further, R (in the case we are imagining) will have the function to carry information about F s, since what would make (i) true in this case is the fact that F is a magnitude with the determinant value v , and for any value v , and margin of error $m \geq 0$, $|v - v| \leq m$. Because of this, (IT') can be understood as a special case of (IT) —namely, the case in which $F = G$. So, in effect, we are able to recover all the advantages of (IT) from the current theory. More generally, any theory τ that amends itself into a theory τ' using (C1) and (C2) should be an extension of τ in the sense that τ is recoverable when the properties F and G are equivalent to one another.

Let us work through the circle case (Case 2 above) to see the theory at work. Plugging into the left hand side of (IT') we have:

(12) A perceptual state R_C represents elliptical surfaces as exemplifying the property *having an aspect ratio of 1 (ellipse(1))* relative to the property *having an aspect ratio of .87 (ellipse(.87))*.

Now, we want to verify that the conditions on the right hand side of the “if” operator to show that such conditions can explain why R_C has the content it has (rather than some other content).

It will be helpful to first verify condition (ii,a). So, the first thing to determine is whether the two properties featured in (12) belong to the same magnitude class, and are approximations of one another. The two magnitudes belong to the magnitude class of *aspect ratio*. Now, recall that for this case in particular the margin of error relevant for determining whether the two properties approximate each other is .15 for the reasons stated above. And, since $|1 - .87| \leq .15$, our two properties do approximate each other.

To verify (i), we want to know whether the following is true:

(13) It is the function of R_C to carry information about elliptical surfaces that approximate the property *ellipse(1)*.

This function ascription is initially plausible since R_C is positively nomically correlated with ellipses with an aspect ratio of .87. Furthermore, it is also plausible given that approximating distal features is sufficient for many tasks perceptual systems engage in, such as navigation and other forms of action guidance (cf. Milner and Goodale 2006, p.106ff.). So, there is nothing perverse in maintaining that perceptual states like R_C aid in guiding action by fulfilling their functions despite only having functions to approximate the features in the distal environment. Moreover, this in turn would explain the utility they provide to their bearer. One could try to deny these latter claims, but it would require them to support a very strong claim, namely that only a perfectly accurate (or near perfectly accurate) representation of the distal environment can be used to guide our actions.

Now, in order to see (C2) at work, we need an explanation for why G s are nomically correlated with perceptual states that represent F . That is, we need to fill in parameter E . For both kinds of cases we have been looking at, we can do so by acknowledging that the visual system is an inferential device of sorts. As many perceptual scientists agree, the key to understanding perception is to understand how perceptual systems are able to construct a representation of the distal environment on the basis of an impoverished stimulus. Thus, it is generally accepted that perceptual systems infer the properties of a distal stimulus σ_d from a proximal stimulus σ_p , induced by σ_d , that underdetermines the properties of σ_d .

It is standard among perceptual scientists to take it that perceptual systems solve this underdetermination problem by making use of the statistical regularities in environment (see, among others, Marr 1982, Howe and Purves 2005; see Burge 2010 for philosophical discussion). However, and this is the important part for us, it is also possible to appeal to such regularities to explain why perceptual systems are subject to systematic error (Howe and Purves 2005).

Such explanations have been fruitfully applied to the kinds of misperception that were discussed in the previous section (as we will see in more detail below). Those cases dealt with visual perception, where the relevant underdetermination problem is going from a 2D retinal image to a 3D representation of what caused that image. So, the claim is that the visual system makes use of the statistical regularities in the environment to predict the most likely cause C of the retinal image when attempting to form a representation of C , and that such regularities sometimes lead the visual system systematically astray. To put things differently, there is a one-to-many mapping between a retinal image I and a 3D cause of I —as is true for any projection from a 3D plane to a 2D plane. So, the visual system has to estimate this cause. Further, it does so by taking into account the statistical regularities of the environment, which sometimes leads to an estimate of the properties of the relevant cause that departs from the actual properties of that cause. We can thus take E for Case 2 to be equivalent to the following proposition:




Figure 6 From Howe and Purves (2002, p.13187): “Physical intervals on the ground plane (colored lines) that extend further and further away from the image plane will project more and more vertically (see corresponding colored lines).”

- (14) Given the statistical regularities in the environment, the most likely distal cause of the proximal stimulus induced by *ellipse(.87)* (which in turn causes R_C) are individuals which possess the property *ellipse(1)*.

An affirmation of (14) requires some empirical evidence. Luckily, Howe and Purves (2002, 2005) have provided strong reason to think that (14) does hold. The argument proceeds from a general investigation of the fact that perceived line length is a function of orientation, which Howe and Purves take to be the cause of our misperception of ellipses. To fix ideas, it will be helpful to see how this argument goes. To do this, it will first be necessary to distinguish between two ways in which a line can be oriented, viz., a line can be oriented frontally, such that a line with a frontal orientation of 90° would be oriented vertically, or a line can be oriented in depth, such that a line with depth orientation of 45° would be tilted backwards at a 45° angle. A line with a depth orientation of 0° would be oriented in the fronto-parallel plane — the plane perpendicular to the subjects line of sight.

Now, the greater a physical line, or interval more generally, is oriented in depth, the smaller the interval appears in the image that it projects on to a 2D plane. That is, the ratio (λ) between the length of a 3D physical interval (L) and the corresponding interval (l) in the image it would project on to a 2D plane increases as the depth-orientation of L (ϕ_L) increases. So, if $\frac{L}{l} = \lambda = 1$ when $\phi_L = 0^\circ$, then $\lambda > 1$, when $\phi_L > 0^\circ$.

Howe and Purves make the following observation:

The more a physical interval extends in depth in the ground plane, the more its projection will tend toward the vertical axis ([Figure 6]). This geometrical fact, coupled with the prevalence of the ground plane, would cause the physical sources of vertical intervals, on average, to incline more in depth and, thus, to be longer compared with the [physical] sources of [2D images of] horizontal intervals. (Howe and Purves 2002, p.13187)

To summarize, Howe and Purves are arguing that if vertical lines are more likely to be oriented in depth, then it will be more likely than not that vertical intervals projected onto the retina are produced by a physical interval larger than the corresponding interval in the projected image. And

that, given the prevalence of the ground plane in our environment and the geometrical fact that the more a line is oriented in depth, the more its projection will tend toward the vertical axis, it is more likely than not that vertical intervals projected onto the retina (a 2D plane) were projected onto the retina by a physical interval larger than the corresponding interval in the image. To verify these claims, Howe and Purves use a technique known as natural scene statistics, the details of which do not concern us here.

So, roughly speaking, if the visual system is inferring the aspect ratio of a shape presented to it in the frontal plane, such as an ellipse, from a 2D retinal image, it will overestimate the vertical axis (relative to the horizontal axis) since it is more likely than not that the vertical axis was projected by an interval oriented in depth, and so more likely than not that the physical interval responsible for the projection is larger than the corresponding interval that was projected onto the retina. In this case, the visual system would get things wrong, but only because it was making a best guess, so to speak, based on the statistical regularities in the environment. Such regularities in this case have lead it astray.

(IT') is well-suited to handle Case 2. Further, we could have chosen to focus on our misperception of angles. Howe and Purves (2005, ch.4) have also provided an explanation for why we misperceive angles in the way we do that is analagous to their explanation of our misperception of aspect ratio. It is of course an open question of how far the present analysis extends. What I take myself to have shown thus far is that the present view at least provides a way for some otherwise plausible theories to make sense of the cases of reliable misperception discussed in §2.2.

5 | OBJECTIONS AND REPLIES

Objection. When E is cashed out for the cases of reliable misperception that we have been discussing it makes (at least an implicit) appeal to the notion of inference. But is that notion not semantic or representational?

Reply. For the sake of argument, let us say 'yes'. But this answer just opens up another line of inquiry, viz.: how we should understand the claim that a perceptual system — all of which I take to be a part of the brain — is making an inference from an underdetermined proximal stimulus (σ_p) to an estimation of its distal cause (σ_d)? To fix ideas, let us suppose, along with many perceptual psychologists (see among others, Knill and Richards 1996), that perceptual systems achieve this task by computing a continuous form of Bayes' theorem:

$$(15) f(\sigma_{d_i} | \sigma_p) = \frac{f(\sigma_p | \sigma_{d_i})f(\sigma_{d_i})}{f(\sigma_p)} = \frac{f(\sigma_p | \sigma_{d_i})f(\sigma_{d_i})}{\int f(\sigma_p | \sigma_{d_i})f(\sigma_{d_i})d(\sigma_{d_i})}$$

$$(16) \max(f(\sigma_{d_i} | \sigma_p))$$

The details of these equations do not matter too much. But a few things are worth noting. First, σ_{d_i} is a continuous variable that ranges over a series of possible distal causes of $\sigma_{d_1}, \dots, \sigma_{d_n}$. Bayes' theorem, in this form, gives us the probability of each member of the sequence given the proximal stimulus, thereby generating a distribution of probabilities. And, (16) is a rule that says to select the σ_{d_i} with the highest probability. Intuitively, this is the best estimate of σ_d given σ_p .

To answer our question, I think we can help ourselves to a classical computationalist understanding of what it is for the brain to carry out such a process. Hohwy, for example, claims that "we can only understand how brains engage in probabilistic inference if we understand how neurons can realize the functional roles set out by forms of Bayes' [theorem]" (Hohwy, p.24). On this




Figure 7 A computational model of Bayes' theorem. g_R is the realization function. And, P_1, P_2, P_3 , and P_4 are physical (neural) states. Compare Egan's (2020, p.27, fig. 1.1.) computational model of an adder.

way of understanding things when we say that the visual system, for example, is computing a certain instance of Bayes' theorem in such-and-such circumstance (such as when a circular surface stimulates the retina) with such and such numerical values, we are also committed to the idea that there is some algorithm that carries out that computation (as per Marr's (1982) second level of analysis) and that that algorithm is in turn realized by brain states. Given this realization relation, we have it that the fact that the visual system is computing Bayes' theorem (or any other statistico-inferential process) is grounded in facts about brain states, namely, that a certain set of brain states in human subjects realizes the functional roles determined by Bayes' theorem, as Hohwy might put it.

In Figure 7, we can see what I have in mind more precisely. $f(\sigma_p|\sigma_{d_i}), f(\sigma_{d_i}), f(\sigma_p)$ are the values needed to compute Bayes' theorem, and thus are understood as the functional, or computational states, needed in order for the system to compute Bayes' theorem. There is then a realization function g_R that specifies the physical states which realize those states. The arrows that go from left to right depict causal relationships between states. When the system is in the physical states that realize the computational states it transitions into the product of those states. I want to emphasize that $f(\sigma_p|\sigma_{d_i}), f(\sigma_{d_i}), f(\sigma_p)$ are equivalent to numerical values. There are a variety of ways these values may be determined. For example, they may be determined by the statistical regularities of the environment (Geisler et al. 2000; Geisler and Kersten 2002 p.509, fig.1), or by the physical properties of the distal stimulus (Geisler and Kersten 2002, p.509, fig.1), or perhaps in some other way (Feldman 2015, p.1016ff.).

Objection. The notion of approximation at play in the account ((A) reprinted below) fails to capture the precise way in which the perceptual states we have been considering represent magnitudes with values that approximate the value of the magnitudes they are nomically correlated with.

(A) A value v_1 approximates a value v_2 relative to a margin of error $m \geq 0$ if and only if $|v_1 - v_2| \leq m$.

For example, since (A) is defined in terms of a margin of error, and the margin of error settled on for the angle case discussed above is 3° , it misleadingly suggests that a perceptual state that represents angles as being 30° is nomically correlated with any angle between 27° and 33° . More generally, it is unlikely that (A) will ever be able to finely capture facts about how a perceptual states content approximates the magnitude(s) with which it is nomically correlated.

Reply. One way to resolve this problem is to introduce an individual approximation metric for each individual case. One can then motivate the choice of metric by appeal to the relevant empirical facts concerning that case. For example, in the angle case there is a curve (Figure 8) that characterizes the degree of misperception for angles between 0 and 180 degrees. Thus, there will

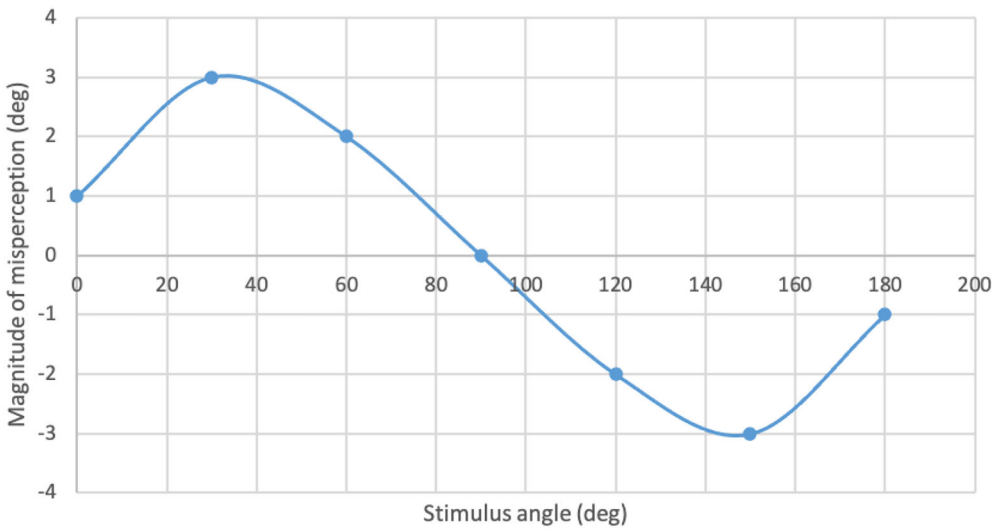


Figure 8 A curve that characterizes the data compiled by Nundy et al. (2000). Cp. Figure 2. The equation that characterizes the curve is roughly $y = 3\sin((2x) + .6\sin(2x))$.

be some equation $y = f(x)$ that characterizes that curve, where the x values correspond to angle sizes and y values correspond to the degree of misperception. Given this equation, we can then say that, for any perceptual state R that represents individuals as being an angle of degree v_r , and (physical) angle v_p that is nomically correlated with R , v_r approximates v_p if and only if $v_r - v_p = y = f(v_p)$. The view here predicts that something similar can be done in other cases of reliable misperception. So, on this revised view of approximation, approximation metrics are determined empirically on a case by case basis.

Objection. The proposal only has a limited appeal, because it can only handle cases of reliable misperception where the relevant perceptual state is plausibly approximately tracking what it purports to represent. But, there is no a priori reason why reliable misperceptions must be limited in this way (see, e.g., Mendelovici 2013).

Reply. There are two things to say here. One is just that the extension that I have proposed is in the spirit of providing an explanatory sufficient condition for determining what a perceptual state represents. Because of this, there is always the possibility that the view will fail to handle various cases. However, and here is the second point, I have focused on *actual* cases of reliable misperception, so there is reason to think that the view provides an explanatory sufficient condition at least for cases that can be found in the actual world. That said however, if one likes they can strengthen (IT') to

(IT'') A perceptual state R represents an individual as being F relative to a property G if and only if

- (i) it is the function of R to carry information about individuals which approximate F s;
- (ii) R is nomically correlated with individuals that exemplify G , where both F and G are of the same magnitude-type \mathcal{M} and the determinate value of G approximates the determinate value of F ; and
- (iii) E , where E explains why G s are nomically correlated with perceptual states that represent F .

And claim to be giving an a posteriori reduction of the *represents* relation (cf. Papineau 2001; cf. Schulte 2020). Such a view would only be held hostage to actual cases, and it would hold with respect to such cases, at least for the time being.

There might be an advantage of limiting the view to an explanatory sufficient condition, however. For example, what may be labeled as “secondary qualities” (such as color), may need a wholly separate treatment from what may be labeled as “primary qualities” (as is suggested by Mendelovici (2013)). If we opt in for (IT') over (IT''), we leave open the possibility that another set of explanatory sufficient conditions will be needed to handle other kinds of cases. Thus, at minimum, we may read the theories gotten from (C1) and (C2) as providing an explanatory sufficient condition for perceptual states that represent primary qualities.

Objection. The cases focused on in §2.2 should not be categorized as reliable misperceptions because they are a part of a system of representations, namely a system of spatial representations, which are by and large truth tracking in the sense that they always get something right about the properties they purport to represent (e.g., the magnitude class to which they belong). Legitimate cases of reliable misperception are cases in which a whole system of representation fails to get anything right about the properties it represents individuals as having. A case like the latter would occur, for example, if color eliminativism turned out to be true (the view that nothing in the world is colored). In that case a whole system of representations, the system of color representations, not only represents inaccurately in the same way all the time, but does not get anything right about the properties it purports to represent (Mendelovici 2016, p.70, fn.9.; cf. Mendelovici 2013, p.424).

Reply. This objection generates a verbal dispute. Whether the cases in §2.2 count as cases of reliable misperception by some pre-theoretic standard is not relevant in the present context. What is relevant is whether the cases I have labeled as reliable misperception (in virtue of the fact that they satisfy the conditions of (RM)) are problematic for theories of perceptual content, and I have argued that they are.

Objection. The cases of reliable misperception discussed in §2.2 are only problematic for theories that privilege the sort atomism assumed in §3. Theories of perceptual content which do not make that assumption will not face the kinds of challenges presented by the cases discussed in §3.

Reply. This objection fails to take into account the possibility that non-atomistic views should be rejected on independent grounds. But even granted that they are legitimate competitors to atomistic theories, it is not clear that they would be compatible with the cases I have been discussing.

One class of theories that is particularly relevant here are consumer based teleosemantic accounts (Millikan 1984, Papineau 1987, Artiga 2013), as it has been argued that such views can account for persistent illusions (like the Ebbinghaus illusion or the Müller-Lyer illusion) which are a lot like the cases discussed in §2.2 (Artiga 2013, p.270f., Mendelovici 2016, p.70, fn.9).²⁰ As already mentioned, a consumer based teleosemantics rejects atomism. Such views assign content to perceptual states that represent whole states-of-affairs, instead of merely assigning contents to perceptual states which come together to generate a representation of a whole state-of-affairs. Artiga (2013) describes a consumer based teleosemantics as follows: “R represents the state-of-affairs P that the consumer system has historically needed in order to perform its function successfully” (2013, p.267). The notion of function relevant for Artiga is that of an etiological normal-proper function (the details of which do not concern us here). I take (CB) (below) to get at what Artiga is saying here in a way that is congenial to human level perceptual

²⁰ Strictly speaking whether or not these arguments go through is not relevant for present purposes because I will argue that a consumer-based teleosemantics fails to capture the cases we have been interested in.

representations, where the cognitive agencies referenced in (CB) are a special kind of consumer, and R is a state-of-affairs (of arbitrary complexity) (adapted from Hill 2010, p.150):

(CB) A perceptual state R represents A if R is assigned to A by a function f such that, for any perceptual state U , $f(U)$ is a state-of-affairs that must hold in order for the cognitive agencies that make use of U to perform their normal-proper functions.

I do not think (CB) fares any better than the other views that we have examined thus far.

Take the 30° angle case, for example. So, there is a perceptual state R_{33° that represents angles as being 33° but is positively nomically correlated with 30° angles and not 33° angles. Suppose (CB) is true for R_{33° and thus a 33° angle state-of-affairs is mapped to R_{33° by $f(U)$ and thus R_{33° must be used by some cognitive agency that needs a 33° angles state-of-affairs to hold in order to perform its normal-proper function.

Now, it is an open question what the relevant cognitive agencies are. Let us for the time being suppose that the cognitive agencies are angle categorizers. One normal-proper function of an angle categorization mechanism that would require $f(U)$ to map the perceptual state R_{33° to a 33° angle state-of-affairs would be a normal-proper function to categorize 33° angles. Such a function would be fulfilled only if the categorization mechanism categorized 33° angles correctly. But, such a function would never go fulfilled if R_{33° is the relevant perceptual state that is used by the categorization mechanism to perform that function, given the lack of positive nomic correlation between R_{33° and 33° angles. This gives rise to the same problems faced by Schellenberg's view and the informational teleosemantics discussed in §3. Namely, the view needs to explain why we should think some mechanism has a normal-proper function that never gets fulfilled.

I think the point generalizes to any possible cognitive agency that may use R_{33° . According to a consumer based account, R_{33° represents 33° angles because that is the state-of-affairs that the consumer system of R_{33° has historically needed in order to perform its function successfully. But since R_{33° is a reliable misperception, any system which used R_{33° to fulfill its etiological normal-proper function and needed a 33° angle state-of-affairs to obtain in order for that function to be fulfilled, would always fail to fulfill its function, since R_{33° is always tokened by something which is not a 33° angle state-of-affairs, and *a fortiori* not the state-of-affairs needed by the cognitive agencies which make use of R_{33° .

Artiga (2013) in particular may respond by saying that “[t]his state [R_{33°], which reliably mis[perceives] is [...] a by-product of [a] representational system that has earned its keep in evolution” (p.271) because “most of the time it produces the right representations” (p.270). I think this response fails to acknowledge that if R_{33° is supposed to be a type of perceptual state with a particular content, then there must be a story of how that content is determined. And that the consumer based account favored by Artiga does not mention that, or tell us how, a state may get its content in virtue of being a by-product of a representational system that has earned its keep in evolution. Accidental by-products are nowhere mentioned in (CB) or in Artiga's informal gloss of (CB). Thus more work needs to be done here. Namely, the official story of the consumer based account needs to be extended to tell us a story of how states can get contents in virtue of being by-products of systems that *most* of the time produce the right representation. Indeed, this may not be possible for the angle case, if the relevant system is a system responsible for angle or shape perception, because as we have seen, both systems seem to get things wrong most of the time (cf. Green and Rabin 2019). Luckily, a modified consumer based account using (C1) and (C2) is presented in the appendix.

Appendix A

Here I briefly show how Fodor's asymmetric dependency theory, Schellenberg's capacitism, and a consumer based teleosemantics can be extended using (C1) and (C2) discussed in §4.

Fodor:

- (F') A perceptual state R represents individuals as being F relative to a property G if
- (i) there is a nomic relation between the property of being a G and the property of being a cause of R tokens, where G and F are of the same magnitude class \mathcal{M} and the determinate value of F approximates the determinate value of G ;
 - (ii) if there are nomic relations between other properties and the property of being a cause of R tokens, then the latter nomic relations depend asymmetrically upon the nomic relation cited in (i); and
 - (iii) E , where E explains why G s are nomically correlated with perceptual states that represent F .

Schellenberg:

- (S') A perceptual state R represents individuals as being F relative to a perceptual capacity C_R , subject S and a property G if
- (i) S 's employment of C_R puts S in R ;
 - (ii) C_R has the function to single out and discriminate properties which approximate F s;
 - (iii) C_R is normally employed in the presence of individuals which exemplify G , where both F and G are of the same magnitude-type \mathcal{M} and the determinate value of F approximates the determinate value of G ; and
 - (iv) E , where E explains why G s are nomically correlated with perceptual states that represent F .

If we opt in for (S'), the possession conditions for C_R would need to be changed to accommodate the fact that its function is to single out and discriminate properties which approximate F s.

Consumer-based teleosemantics, where A and B are state-of-affair types of arbitrary complexity:

- (CB') A perceptual state R represents A relative to some B if
- (i) the values of the magnitudes that are featured in A approximate those featured in B ;
 - (ii) B is assigned to R by a function f such that, for any perceptual state U , $f(U)$ is a state-of-affair type whose holding is sufficient for the cognitive agencies that make use of U to perform their normal-proper functions; and
 - (iii) E , where E explains why states-of-affairs of type B are nomically correlated with perceptual states that represent states-of-affairs of type A .

ACKNOWLEDGEMENTS

I would like to thank Frances Egan, Seth Goldwasser, Paul Pietroski, Susanna Schellenberg, Daniel Stoljar, and an anonymous referee at this journal for written comments on earlier versions of this paper. I also owe much to discussion with EJ Green, Brian McLaughlin, and Ezra Rubenstein. I am further indebted to audiences at the Australian National University (twice, sorry!), Australian Catholic University, Monash University, Ruhr University Bochum, Rutgers University,

and the 2022 Joint Session of the Society for Philosophy of Psychology and the European Society for Philosophy of Psychology.

REFERENCES

- Artiga, Marc. 2013. Reliable Misrepresentation and Teleosemantics. *Disputatio*, 37(37), 265–281.
- Brewer, Bill. 2011. *Perception and It's Objects*. Oxford: Oxford University Press.
- Burge, Tyler. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.
- Carpenter, R.H.S., & Blakemore, Colin. 1973. Interactions between Orientation in Human Vision. *Experimental Brain Research*, 18, 287–303.
- Craven, B.J. 1993. Orientation Dependence of Human Line-Length Judgements Matches Statistical Structure. *Proceedings: Biological Sciences*, 253(1336), 101–106.
- Dretske, Fred I. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Dretske, Fred I. 1988. *Explaining Behavior*. Cambridge, MA: MIT Press.
- Egan, Frances. 2020. A Deflationary Account of Mental Representation. In: Joulia Smortchkova, Krzysztof Dolega, & Schlicht, Tobias (eds), *Mental Representations*. New York: Oxford University Press.
- Feldman, Jacob. 2015. Bayesian Models of Perceptual Organization. *Pages 1008–1026 of: Wagnmans, Johan (ed), The Oxford Handbook of Perceptual Organization*. Oxford: Oxford University Press.
- Fodor, Jerry. 1987. *Psychosemantics*. Cambridge, MA: MIT Press.
- Fodor, Jerry. 1990. *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.
- Ganson, Todd. 2021. An Alternative to the Causal Theory of Perception. *Australasian Journal of Philosophy*, 99(4), 683–695.
- Geisler, Wilson S., & Kersten, Daniel. 2002. Illusions, Perception and Bayes. *Nature Neuroscience*, 5(6), 508–510.
- Geisler, W.S., Perry, J.S., Super, B.J., & Gallogly, D.P. 2000. Edge Co-occurrence in Natural Images Predicts Contour Grouping Performance. *Vision Research*, 41(6), 711–724.
- Godfrey-Smith, Peter. 1991. Signal, Detection, Action. *Journal of Philosophy*, 88(12), 709–722.
- Green, E.J., & Rabin, Gabriel Oak. 2019. Use Your Illusion: Spatial Functionalism, Vision Science, and the Case Against Global Skepticism. *Analytic Philosophy*, 0(0), 1–34.
- Hibbard, Paul B., Goutcher, Ross, O’Kane, Lisa M., & Scarfe, Peter. 2012. Misperception of Aspect Ratio in Binocularly Viewed Surfaces. *Vision Research*, 70, 34–43.
- Hill, Christopher. 2010. Intentionality Downsized. *Philosophical Perspectives*, 20, 144–169.
- Howe, Catherine Q., & Purves, Dale. 2002. Range Image Statistics can Explain the Anomalous Perception of Length. *Proceedings of the National Academy of Sciences of the United States of America*, 99, 13184–13188.
- Howe, Catherine Q., & Purves, Dale. 2005. *Perceiving Geometry: Geometrical Illusions Explained by Natural Scene Statistics*. Springer Verlag.
- Howhy, Jakob. 2013. *The Predictive Mind*. Oxford: Oxford University Press.
- Knill, David C., & Richards, Whitman. 1996. *Perception as Bayesian inference*. Cambridge University Press.
- Marr, David. 1982. *Vision*. San Francisco: W.H. Freeman.
- Matthen, Mohan. 1988. Biological Functions and Perceptual Content. *Journal of Philosophy*, 85(1), 5–27.
- McLaughlin, Brian P. 2016. The Skewed View from Here: Normal Geometrical Misperception. *Philosophical Topics*, 44(2), 231–299.
- Mendelovici, Angela. 2013. Reliable Misrepresentation and Tracking Theories of Mental Representation. *Philosophical Studies*, 165(2), 421–443.
- Mendelovici, Angela. 2016. Why Tracking Theories Should Allow for Clean Cases of Reliable Misrepresentation. *Disputatio*, 8(42), 57–92.
- Mendelovici, Angela, & Bourget, David. 2014. Naturalizing Intentionality: Tracking Theories versus Phenomenal Intentionality Theories. *Philosophy Compass*, 9(5), 325–337.
- Millikan, Ruth Garrett. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Milner, A. David, & Goodale, Melvyn A. 2006. *The Visual System in Action*. Second edn. Oxford: Oxford University Press.
- Neander, Karen. 1995. Misrepresenting Malfunctioning. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 79(2), 109–141.

- Neander, Karen. 2017. *A Mark of the Mental: In Defense of Informational Teleosemantics*. Cambridge, MA: MIT Press.
- Nundy, Surajit, Lotto, Beau, Coppola, David, Shimpi, Amita, & Purves, Dale. 2000. Why are Angles Misperceived. *Proceedings of the National Academy of Sciences of the United States of America*, 97(10), 5592–5597.
- Papineau, David. 1987. *Reality and Representation*. New York: Blackwell.
- Papineau, David. 2001. The status of teleosemantics, or how to stop worrying about Swampman. *Australasian Journal of Philosophy*, 79, 279–289.
- Peacocke, Christopher. 1989. Perceptual Content. In: Almog, J., Perry, John, & Wettstein, Howard K. (eds), *Themes from Kaplan*. Oxford: Oxford University Press.
- Peacocke, Christopher. 2015. Magnitudes: Metaphysics, Explanation, and Perception. *Pages 357–387 of: D. Moyal-Shaprock, V. Munz, & Colvir, A. (eds), Proceedings of the 36th International Symposium*. Berlin: de Gruyter.
- Price, Carolyn. 2001. Determinate Functions. *Nous*, 32(1), 54–75.
- Rubner, Andrew. 2022. Normal-proper functions in the philosophy of mind. *Philosophy Compass*, 17(7), 1–11.
- Schellenberg, Susanna. 2014. The Relational and Representational Character of Perceptual Experience. In: Brogaard, Berit (ed), *Mental Representations*. New York: Oxford University Press.
- Schellenberg, Susanna. 2018. *The Unity of Perception: Content, Consciousness, Evidence*. Oxford: Oxford University Press.
- Schellenberg, Susanna. 2020. Capacities First. *Philosophy and Phenomenological Research*, 100, 744–757.
- Schulte, Peter. 2018. Perceiving the World Outside: How to Solve the Distality Problem for Informational Teleosemantics. *Philosophical Quarterly*, 68(271), 349–369.
- Schulte, Peter. 2020. Why mental content is not like water: reconsidering the reductive claims of teleosemantics. *Synthese*, 197, 2271–2290.
- Shea, Nicholas. 2018. *Representation in Cognitive Science*. Oxford: Oxford University Press.
- Skyrms, Brian. 2010. *Signals: Evolution, Learning, and Information*. Oxford: Oxford University Press.
- Zhu, Jlelei Emma, & Ma, Wei Ji. 2020. Orientation-dependent Biases in Length Judgments of Isolated Stimuli. *Journal of Vision*, 17(2), 1–19.

How to cite this article: Rubner, A. (2024). Theories of perceptual content and cases of reliable spatial misperception. *Philosophy and Phenomenological Research*, 108, 430–455. <https://doi.org/10.1111/phpr.12980>