Reconsidering Perceptual Content*

William T. Wojtach†‡

An important class of teleological theories cannot explain the representational content of visual states because they fail to address the relationship between the world, projected retinal stimuli, and perception. A different approach for achieving a naturalized theory of visual content is offered that rejects the traditional internalism/externalism debate in favor of what is termed "empirical externalism." This position maintains that, while teleological considerations can underwrite a broad understanding of representation, the content of visual representation can only be determined empirically according to accumulated past experience. A corollary is that a longstanding problem concerning the indeterminacy of visual content is dissolved.

1. Introduction: The Underdetermination of Retinal Stimuli. In "An Essay Towards a New Theory of Vision" ([1709] 1975), Berkeley posed a puzzle that theorists have had to grapple with ever since. Given the transformations that occur when sources[1] in three-dimensional space project onto a two-dimensional surface, the retinal image cannot be used to determine the size, distance, and orientation of the real-world geometry that produced the stimulus, since the full dimensionality of the world is not preserved. This problem becomes only more intractable when it is recognized that the quality and quantity of light returned to the eye further intertwines the relative contributions of illumination, surface reflectance, and atmospheric transmittance, thereby making it impossible to recover how each of these factors have been combined to produce the patterns of light

*Received September 2006; revised November 2008.

†To contact the author, please write to: Department of Philosophy, Center for Cognitive Neuroscience, Box 90999 LSRC, Duke University, Durham, NC 27708; e-mail: wtw3@duke.edu.

‡I would like to thank Güven Güzeldere, Alexander Rosenberg, Dale Purves, and Fred Dretske for helpful comments on earlier drafts of the manuscript.

1. The term 'source' is used instead of 'object' in an effort to remain neutral regarding the conditions in the world that give rise to projected stimuli. Also, 'image' and 'stimulus' are equivalent terms in what follows, both referring to the patterns of light projected on the retinal surface.

Philosophy of Science, 76 (January 2009) pp. 22–43. 0031-8248/2009/7601-0001$10.00

projected on the retina. As a result of such circumstances, retinal stimuli are underdetermined with respect to the world.

Stated in such terms, the fact that projected images are underdetermined might not seem to be a significant problem since we do not perceive the retinal image. But if retinal stimuli contain the only direct visual information an observer has about the world, yet such stimuli cannot uniquely specify the conditions of that world, then explaining how observers generate visual representations that are useful guides to the *sources* of stimuli presents a profound challenge (cf. Purves and Lotto 2003; Wojtach 2005). This puzzle has come to be known as the "inverse optics problem" (Palmer 1999).[2]

To better understand the geometric aspects of the inverse optics problem, consider the relationship between the projection of a straight line on an image plane (or retina) and the real-world conditions that could have generated the image (Figure 1). From the projected stimulus alone, it would be impossible to determine the three-dimensional source that is the unique cause of the projection, since any number of sources of different sizes, at different distances from the observer, and at different orientations in space could all subtend the same visual angle in the image. Notice that this problem remains regardless of whether the image is a static scene, as in Figure 1, or a dynamic scene, as in Figure 2.

If the sources of Figure 1 were moving in three-dimensional space from right to left (Figure 2, frames A to C), then the projected image sequence

2. Some theorists, most notably Gibson (1979) and those sympathetic with his view, do not accept the implications of the inverse problem as presented here. While a full treatment of this debate is outside of the scope of this article, two points should be mentioned. First, a common misunderstanding is that whatever significance the inverse optics problem has, it pertains only to impoverished stimuli. If the stimuli were more complex, it is believed, then the problem of underdetermination would effectively disappear, since the features of such stimuli could then be used to uniquely represent the conditions in the world. Although this may seem correct, it should be clear that the inverse problem remains regardless of whether retinal stimuli originate from natural scenes or are derived from laboratory settings. Because there is nothing inherent to the stimuli produced from natural scenes that is lacking in stimuli produced from artificial scenes, neither class of stimuli are immune to the transformations that occur when three-dimensional space is projected onto the retina; nor do they differ with respect to the conflation of illumination, reflectance, and transmittance in the image. Thus, relying on "texture gradients" and other aspects of natural (and artificial) scenes in an attempt to overturn—or worse, ignore—the inverse problem is misleading. Another common misunderstanding is that, even if the inverse optics problem does exist it only pertains to static images, since the information present in moving images would permit a one-to-one relationship with the physical source. As discussed in the text, however, even in the case of motion the inverse optics problem is not diminished, since the extra dimension of time cannot resolve the underdetermination that exists between images and their generative sources.

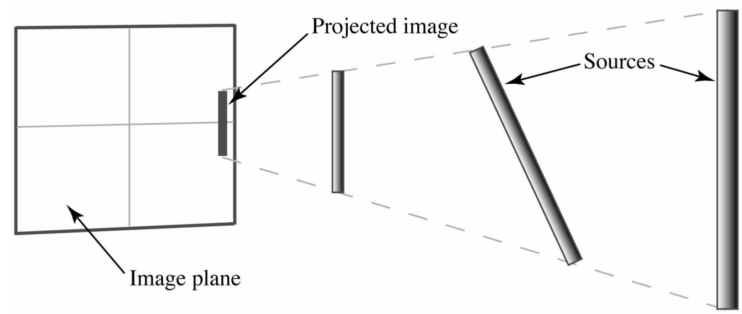

Figure 1. The inverse optics problem. Due to perspective transformation, objects of different sizes, at different distances, and in different orientations can all generate the same projection on an image plane. Therefore, a projected image cannot uniquely specify its source in three-dimensional space.

illustrates a further aspect of the inverse problem: in addition to the problems posed by size, distance, and orientation, the *speed* of a source cannot be uniquely determined with respect to a moving image on the retina (Wojtach et al. 2008). Therefore, even though the series depicted in Figure 2 contains only three different sources moving in multiple frontal planes with various speeds, an infinite number of sources of different sizes, at different distances, with various orientations, and traveling at different speeds could have produced the stimulus. As a result, the inverse problem pertains to full range of speeds and geometries that can arise from three-dimensional space.

As mentioned above, the conflation of information also affects the quality and quantity of light projected onto the retina. Perhaps the easiest way to capture this aspect of the inverse problem is to consider achromatic stimuli—in particular, the puzzle introduced by simultaneous brightness contrast (Figure 3).[3] As is well known, luminance is an objective measure of the intensity (or quantity) of light returned to the eye from physical sources as measured by a photometer. We do not perceive luminance, however, but lightness and brightness, typically measured by having an

3. It is important to note that the problems introduced here in terms of achromatic stimuli also extend to chromatic stimuli, as is evident in the well-known problems of simultaneous color contrast and color constancy. The inverse problem therefore affects *all* patterns of light on the retina, not just the class of achromatic stimuli presented here for illustrative purposes.

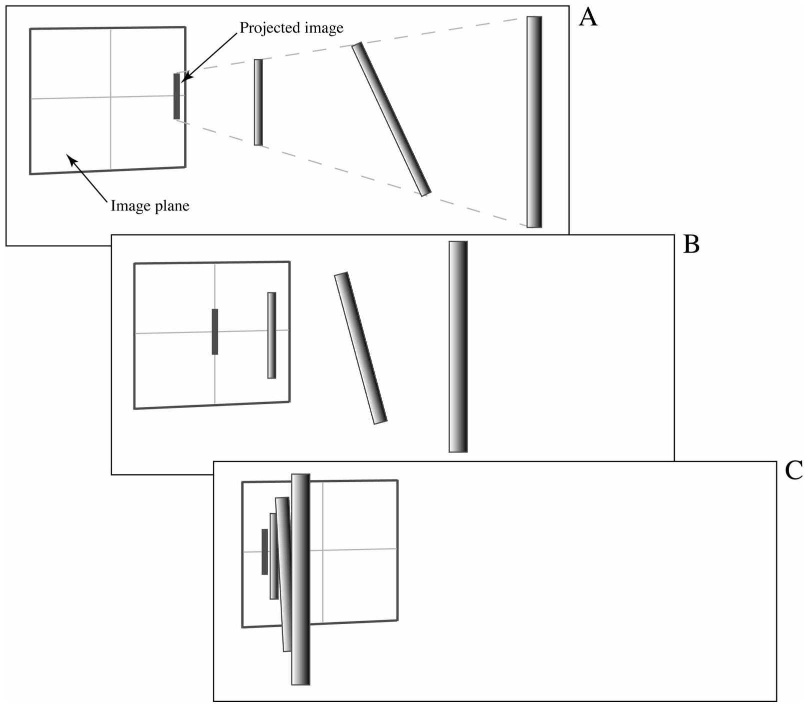

Figure 2. The inverse optics problem as it pertains to moving objects. (A) Three objects of different sizes, at different distances, and in different orientations project the same image, as in Figure 1. (B) The same objects moving with different speeds in space can generate the same sequence of projections on the image plane (see also [C]). Similar to a static image, then, a moving image cannot uniquely specify its source in three-dimensional space.

observer indicate the appearance of one surface relative to another.[4] Intuitively, one might think that if two surfaces in a scene returned the same amount of light to the eye (i.e., had the same luminance), then those surfaces would be perceived as equally bright; as demonstrated by the example of simultaneous brightness contrast, however, this is not the case.

In Figure 3, two small target rectangles of equal luminance are each

4. The terms 'lightness' and 'brightness,' though usually conflated, do refer to different subjective assessments. 'Lightness' concerns the appearance of a given surface due to how much light that surface *reflects* to the eye relative to the other surfaces in the scene. 'Brightness,' on the other hand, pertains to the apparent intensity of light due to the surface in question *emitting* (rather than reflecting) the light returned to the eye. Here we will follow convention, however, and use the term 'brightness' where 'lightness' should technically be employed, since the surfaces in question are only reflecting, not emitting, light.

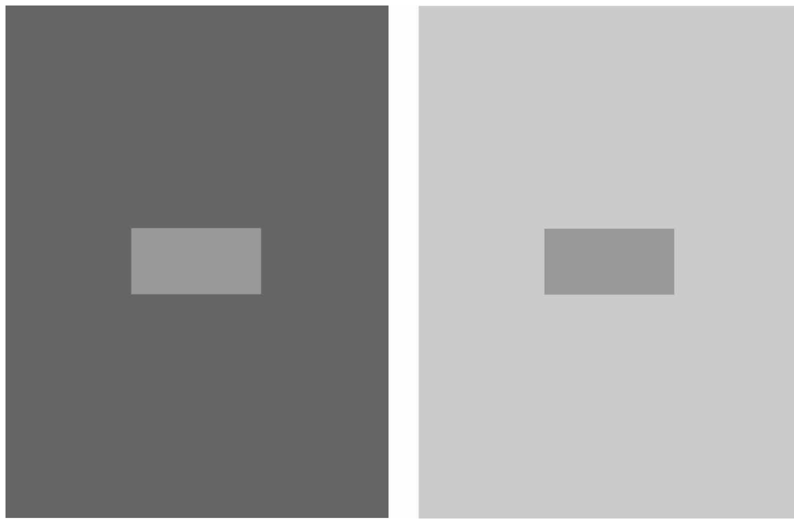

Figure 3. Simultaneous brightness contrast effect. When two target surfaces returning the same physical amount of light to the eye (i.e., having equal luminance) are each surrounded by surfaces returning different physical amounts of light (i.e., having different luminances), the targets appear to be differently bright. In this case, the small target rectangle surrounded by the background of relatively low luminance on the left appears to be brighter than the small target rectangle surrounded by the background of relatively high luminance on the right. See text for the significance of this effect.

surrounded by larger backgrounds with different luminances. When the targets and their respective backgrounds are presented adjacent to one another, so that they can be compared, the targets appear to be *differently* bright: the target surface surrounded by the background of relatively low luminance on the left appears to be brighter than the target surface surrounded by the background of relatively high luminance on the right. Thus, simultaneous brightness contrast is a simple means to illustrate an important point: our visual experience is not a veridical representation of either the retinal image or the underlying reality of sources in the world.

While the physiological details concerning how the visual system generates perceptions of brightness remain enigmatic, the underlying rationale for *why* it does so can be offered by taking into account the circumstances posed by the inverse problem. For unless the visual system had some way of "reverse engineering" luminance values to reveal the actual conditions of illumination, surface reflectance, and atmospheric transmittance that together generated the stimulus, relying on luminance alone (or *any* property of the retinal stimulus as such) to determine the content of visual experience would be a poor method for representing the world and guiding behavior. As described below and in Section 4, there are good reasons to believe that a better tactic—one that results in perceptions of brightness, not luminance—evolved to contend with such circumstances.

Whereas Berkeley argued that the inverse problem could be overcome by coupling visual and tactile information, Helmholtz hypothesized that past experience might be relied upon to supplement the information in the current retinal stimulus (Helmholtz [1866] 1924). Described as a process of "unconscious inference," the use of such experience could generate perceptions more appropriate to the sources that generated the retinal stimulus than might be achieved by using the properties of the stimulus alone. More recently, this general idea that the visual system must in some way use past experience broadly construed[5] to inform perception and behavior has been employed to great advantage with what have come to be called "probabilistic" models, thereby formalizing Helmholtz's proposal of how the brain might effectively contend with the inverse optics problem.[6]

The underlying framework common to all probabilistic models is that evolution by natural selection has shaped the design of perceptual systems in response to environmental conditions. Because of this, the statistical regularities of natural environments can be used to investigate how organisms have come to represent the world (Simoncelli and Olshausen 2001). When adopting a probabilistic approach, however, the chosen method can impact how one understands visual representation and the determination of content. For example, some probabilistic models apply Bayes' theorem to compute the probability distributions of real-world states that could have given rise to a retinal stimulus (cf. Knill and Richards 1996; Maloney 2001; Rao, Olshausen, and Lewicki 2002; Weiss, Simoncelli, and Adelson 2002; Stocker and Simoncelli 2006). While the perceptual "decision" resulting from this technique can be expressed in various ways (e.g., as a maximum a posteriori, or MAP, rule for selecting

5. For any model that relies on past experience, the relevant parameters are both the phylogenetic experience of the species (the process of evolution by natural selection that developed the visual system over time) and ontogenetic experience (individual learning via behavioral feedback). The precise role that phylogeny and ontogeny each have in generating visual representations remains an unresolved issue.

6. It should be mentioned that inverse problems do not only occur in vision. For example, there is a similar problem that exists in audition—what has been called the "inverse acoustics problem"—that arises because a given variation in sound pressure can originate from infinitely many different combinations of the initiating mechanical force, the resonant properties of the body or bodies acted on, and qualities of the intervening medium between the source and the listener. The apparent success of sound localization suggests that the auditory system, like the visual system, has evolved to contend with such circumstances using past experience.

the mode of the posterior distribution), Bayesian models tend to assume that a visual representation corresponds to the most likely physical source of a stimulus. Due to the constraints posed by the inverse optics problem, however, this may not be the best way to conceive of representation.

Another probabilistic model takes a different approach. Given that our visual experience does not conform to the physical properties of the retinal stimulus or the underlying source that produced the stimulus, representation can be conceived in terms that do not correspond to either. On this view, called empirical ranking theory, representations are not generated according to what the most likely properties of the retinal stimulus or real-world source might be, but by the relative rank a given retinal stimulus has in relation to all other instances of the same (or similar) stimulus parameter in accumulated past experience (Yang and Purves 2004; Howe and Purves 2005; Wojtach 2005; Long, Yang, and Purves 2006, Wojtach et al. 2008). As a result, visual representations are taken to be empirical constructs that cannot be directly mapped onto reality.

To better understand this position, consider the puzzle of simultaneous brightness contrast illustrated in Figure 3. Instead of attempting to explain the perception of brightness in terms of the luminance values of the stimulus or the most probable surface reflectances of the source, the representation of differential brightness elicited from the equiluminant target stimuli is determined by how often such targets co-occurred with the full range of surrounding luminance values in natural environments. This strategy therefore relies on more than the luminance values of the target and surround in the present stimulus; instead, it depends on the entire distribution of conjoint luminance values in accumulated past experience. When these conjoint values of target and surround are ordered on a scale ranging from the lowest to the highest luminance according to their frequency of occurrence, a procedure called "histogram equalization" in information theory (Laughlin 1981), the same target stimulus will have a different relative rank (corresponding to a different perceived brightness) depending on the surrounding context. As a result, the representation of brightness generated from the quantity of light falling on the retina can differ markedly, giving rise to the differences between perceived brightness and luminance (or any other visual quality and the measured properties of the stimulus) that are commonly noticed in the case of "visual illusions." If this framework is right, then contending with the underdetermination of retinal images by relying on the relative rank of a given stimulus in accumulated past experience is the key to understanding representation and the determination of content.[7]

7. Further details and implications of the empirical ranking approach are provided in

https://doi.org/10.1086/597020 Published online by Cambridge University Press

2. Teleological Theories and Externalism. One immediate implication of the inverse optics problem is that internalism—the position that the content of visual states can be characterized with reference to internal factors alone[8]—must be incorrect, since external states of affairs will be essential in determining any visual representation. For as indicated by the inverse problem itself, there must be some way to link underdetermined images with external sources in order to generate useful percepts of the world. Although internal states are a necessary condition for visual perception, they cannot be sufficient to account for representation or the determination of visual content.[9] This seems to leave externalism as the only viable candidate for an explanation of visual representation; therefore, further clarifying remarks about internalism will be set aside.[10]

To explain visual representation, two topics are of central concern: misrepresentation and the determination of content. Of the various externalist accounts that have been proposed, the naturalized perspectives offered by Dretske (1981, 1986, 1995), Millikan (1984, 1993), Fodor (1987), and others under the rubric "teleological semantics" constitute perhaps the most significant attempts to address these problems. Despite a number of insights from this perspective, however, the underlying reason why these efforts have been unsatisfactory is because the inverse optics problem and its consequences are largely obscured. For as argued below, once the inverse problem is appreciated, the very idea of a "correct" and "incorrect" representation, as well as how content is determined, must be reconceived. In an effort to demonstrate this, particular attention will be paid to Dretske's teleological framework (1981, 1986, 1995), since his view most easily illustrates these matters; it could be contended, however, that similar problems arise on other teleological accounts as well.

Section 4. For a discussion of the strengths and weaknesses of both Bayesian and empirical ranking models, see Howe, Lotto, and Purves 2006.

8. Although a pure internalist approach is not widely held, and is often replaced by a dual aspect view of "narrow" and "wide" contents (e.g., Block 1986), the weaknesses that are inherent in any account that advocates internalism are important to note once the inverse problem is acknowledged.

9. A welcome result from this perspective is the recognition that fantastic thought experiments like variants of Putnam's (1975) "twin earth" or Davidson's (1987) "swampman" scenario are not required to argue against the role of internalism in visual perception. We need look no farther than what occurs in *this* world and in every day, nonmiraculous swamp circumstances to understand the problems inherent with internalism.

10. The philosophical debates concerning internalism and externalism are extensive, and include intricate arguments offered in defense of both positions. In order to make some progress on these matters, however, it is necessary to suppress many of these points here. The argument offered in the following sections should indicate why this strategy was required.

Dretske ultimately secures the ability of something to represent in the capacity of what he calls "natural signs" to indicate: shadows to the east mean (indicate) that the Sun is in the west; expanding metal means (indicates) that the temperature is rising (Dretske 1986). These signs mean what they do whether or not anyone or anything realizes it, and they can gain their meaning in two ways. One way is through lawful relations between the sign (or the sign's having a certain property) and the condition that constitutes its meaning—e.g., a shadow to the east (the sign) and the Sun being in the west (the condition). Another way to achieve the connection between natural signs and their significance is through nonlawful regularities—i.e., conditions that reflect regularities, but do so without the objective constraints of a law. Dretske's preferred example is that of a ringing doorbell (the sign) and the meaning it carries that someone is at the door (the sign's significance). While not a lawful relation, a regularity has been established such that when the bell rings this means (indicates) that a person, and not something else, is at the door. As long as this regularity persists, a ringing doorbell retains what Dretske calls its "natural meaning" (Dretske 1986, 19–20).[11]

An important limitation of natural signs, however, is that while they can indicate (and thus fail to indicate) something about the world, they cannot *falsely* indicate anything. If the Sun is not in the west, then shadows to the east do not retain their natural meaning that the Sun is in the west; similarly, the ringing doorbell cannot be a natural sign that someone is at the door when no one is there (Dretske 1986, 20–21). So construed, natural signs cannot be said to represent their conditions, since they cannot offer the ability to misrepresent. To explain this ability—and therefore to explain "genuine" representation—Dretske argues that a connection must be made between the functional roles of systems and natural signs. If a system (or component of a system) has an identifiable function, then there is a way to speak of that system (or component) as misrepresenting the conditions it has otherwise been designed to indicate; this is the system's "functionally derived meaning" (Dretske 1986, 22).

To understand Dretske's aim, consider the more problematic case of the ringing doorbell. In normal circumstances, the doorbell system naturally means (indicates) that someone is at the door. Yet the doorbell, as a designed system with an identifiable function, still means something in *abnormal* circumstances; when it is malfunctioning, it can mean that someone is at the door when in fact no one is there. Thus, according to Dretske,

11. "Natural meaning," then, is equivalent to indication. If A naturally means B, then A indicates B. This is just another way of talking about information: there is an informational relationship between A and B—namely, A carries the information that B (Dretske, personal communication; also see Dretske 1981).

the doorbell system can functionally mean (represent) something about the conditions it was designed to track because under *normal* circumstances it has natural meaning (indication).[12] For Dretske, such conditions are constitutive of misrepresentation and, therefore, representation.

This same general point can be made for biological systems: if it is the natural function of a system (or component) to normally indicate the conditions it has evolved to signify, then under abnormal conditions, when the system fails to indicate something about the conditions it has evolved to signify, the system still retains a functionally derived meaning because it was supposed to naturally mean what it evolved to indicate. If this is correct, then we seem to have a teleological basis of the way content (and hence false content) might be instantiated in such systems (Dretske 1986, 22–25).

While a teleological approach seems to be a significant advancement in the effort to explain (mis)representation, the problem of how such representational content is determined remains. A particular strength of Dretske's position is that it offers a way to contend with the indeterminacy of content.[13] The problem here is whether the representational content of a visual state is of an F in the world, or of an F-or-G-or-H-or-. By relying on teleological function this problem is apparently overcome, since the content of a representational state is now determined in terms of a system's natural function—what a system (or component) is *supposed* to indicate. Therefore, if it is the natural function of a system to indicate an F, rather than a G or an H, then that is what the system naturally (hence functionally) means; that is what the system represents. The question is whether such an appeal to teleological considerations can guarantee this result.[14]

12. To emphasize an important point, for Dretske there is no sense in which the doorbell naturally means (indicates) that someone is at the door when no one is there. That is what it nonnaturally (functionally) means; that is what it *represents*. Thus when the topic is functional meaning, the emphasis shifts from indicating to representing—and this, according to Dretske, is the crux of (mis)representation.

13. In this respect, the most widely known criticism concerning the inability of a teleological theory to determine content is by Fodor (1990).

14. While teleological considerations claim to solve this (distal) form of the problem of indeterminacy, Dretske states that we simply need to accept that even sophisticated systems will be plagued with another sort of indeterminacy, since there will always be the possibility of describing functional meaning in terms of a disjunction of proximal input. For largely pragmatic reasons, however, Dretske maintains that this kind of indeterminacy is not a substantial problem, since the point of representation is with regard to the distal source, not the proximal stimuli. For the details of this argument, see Dretske 1986, 26–35. While this approach seems correct—i.e., we do not represent images on the retina, but their generative sources (see Section 1)—the argument in

3. The Importance of the Inverse Optics Problem. Teleological frameworks like Dretske's gain leverage on the problem of representation because they assume that since the function of simple systems can be straightforwardly linked to natural signs, the same holds for the function of complex systems. For reasons that stem from the inverse optics problem, however, this otherwise plausible assumption does not hold in the case of vision. First, consider that while Dretske's concept of a natural sign is unimpeachable, this concept only succeeds because signs are linked to their significance through lawful relations or nonlawful regularities in a one-to-one manner, as when a shadow to the east indicates that the Sun is in the west (as opposed to a light bulb to the north). Notice, however, that in the context of perception and the inverse optics problem, natural signs that work in one direction (from source to projected image)[15] do not work in the other: an image (the sign) cannot indicate the *specific* condition in the world that is its source (the sign's significance), since the mapping in this direction is a one-to-many relation (cf. Figures 1 and 2). This is problematic if visual representation is to be explained by connecting the functional role of the system to natural signs: even if the visual system has the identifiable function of visually representing the *source* of a projected image, the inability to link this function with a natural sign undermines the basis for a teleological explanation of representation.

Suppose, however, that despite these difficulties sense *could* be made of a natural sign in the context of the inverse optics problem. The obstacles to a teleological framework of visual representation still remain, since the link between indication and function is achieved by appealing to "normal" and "abnormal" circumstances. For once the idea that a system has evolved to represent external conditions is in place, it is a small step to the position that under normal conditions it has evolved to represent those conditions *correctly*. With this move now made, a natural corollary is to maintain that when normal conditions in the environment fail, so does the representation. Misrepresentation, then, comes to be identified with the states of systems that are in error because they do not correctly represent the external conditions on which the functionality of the system

Section 4 will question whether Dretske's approach can even solve the more important problem of distal indeterminacy as raised by Fodor (1990).

15. In fact, however, this may be granting too much to the notion of a natural sign. For example, if a particular combination of illumination, surface reflectance, and atmospheric transmittance are to be considered a natural sign of an image with a certain luminance value, then when a different combination of these factors produce the same luminance value, is this to be considered a *different* natural sign? If so, why? Rather than attempting to adjudicate such matters, the important point is that *even if* the concept of a natural sign can function from source to stimulus, it cannot function from stimulus to source (see text).

is predicated. Thus, the implication of a teleological approach is that the functional system in question either operates properly, as it was designed to do, or it is malfunctioning; it is either representing correctly, or it is not.[16]

To make this point clear, consider the following. A teleological framework like Dretske's assumes that because the visual system

1. evolved to (is supposed to) represent external conditions,

the functionally derived meaning of the system can

2. represent those conditions *as they really are*.

But 2 contains a further commitment, namely,

3. what the visual system *represents* is what we *see*.[17]

There is much to be said for the position that a biological system's natural function is to represent its environment, as captured in statement 1. One obvious benefit is that now evolutionary history can be used to define the representational parameters of the system (i.e., what it represents), rather than having to consider the visual system's present state and proceed from those facts alone (cf. Millikan 1993). But while such reasoning seems to apply unproblematically to statements like 1, this general notion of representation cannot be extended to include the phenomenology of visual representations, as expressed in statements 2 and

16. Even if one holds that the difficulty here is in elucidating what should count as the abnormal conditions under which misrepresentation takes place (or, conversely, to describe precisely what should count as the normal conditions for when genuine representation occurs), the fundamental assumptions just noted remain. For instance, Millikan (1984, 1993) avoids having to discuss what constitutes normal and abnormal conditions directly in favor of an *indirect* approach: the notion of indicating correctly is best thought of as a biological norm for our perceptual systems. As long as our perceptual systems are operating in such a way that accounts for the survival of the species, a representation does not occur without its represented. While Millikan's target here is slightly different from the goals of the present article, like Dretske, she maintains that it is by first specifying how a biological system is *supposed* to work that we can determine what should count as the normal (and hence abnormal) conditions for genuine representation (and misrepresentation). Although this is a significant insight, it still maintains that our representations are to be clustered into one of two discrete groups: those that are correct and those that are not, with representation and misrepresentation aligned accordingly. Hence, while Dretske and Millikan have both highlighted that we must look to the evolutionary design of the system if we are to explain representation, they still want to maintain that it makes sense to conceive of representation in terms of correctness and incorrectness.

17. In a teleological framework (as well as in most others concerning vision) this might seem tautologous. Certain distinctions will be drawn in the following section to illustrate why statement 3 is problematic as it stands.

3: for given the relationship between an image and its source as described by the inverse problem, the ability to characterize the representational product of vision as "correct" or "incorrect" must be questioned.

This is a contentious claim. It might be argued, therefore, that while the inverse problem poses a challenge for the way organisms visually represent the world, it does not follow that the very notion of a correct (or incorrect) representation as expressed in a teleological framework is called into question. For suppose that all the visual system can do is *contend with* (rather than *solve*) the inverse problem. When it contends with this problem successfully, by generating a representation of the real-world source of the stimulus, then this is all that is required to capture what is at issue with a correct representation; conversely, when a representation is not of the source that generated the stimulus, this would describe an incorrect representation. Could not a teleological approach still operate with these modified notions of correctness and incorrectness in place, and thereby explain (mis)representation with respect to vision?

While this reasoning seems to address the issue, it fails to notice some important consequences ushered in by this shift in the meaning of terms. First, by using 'correct' to now mean "contend with successfully," the term is no longer applied in its original sense. Whereas before it was meant to convey something like veridicality, or representing the world in accord with normal conditions, it now carries a meaning far from the manner in which it was initially employed. A "correct" representation in this revised sense—i.e., where the term refers to some criterion of *accessibility* for representational states—is not what Dretske (and others) had in mind when offering a teleological explanation of (mis)representation or the determination of content.

Suppose, however, that something like this modified notion of a correct representation were the aim of a teleological position, the objective being to secure the conditions for assessing a given representation. Notice that if this were the case, the result is no better. For regardless of how "assessability" is defined in this context, such a position would have to maintain that *the* source of a stimulus could be represented, since only then could the representation be assessed. Indeed, a teleological approach must maintain this, given that the foremost strength of this perspective is the professed ability to explain the difference between a representation and a misrepresentation. As the inverse problem demonstrates, however, visually representing *the* source of a stimulus cannot be accomplished. Therefore, even allowing a representation to be couched in such terms cannot salvage a teleological approach. Without a way to represent (and thereby assess) *the* source, the defining characteristic of a (mis)representation no longer remains.

To evaluate these conclusions in terms of a concrete example, consider the simultaneous brightness contrast effect of Figure 3. As argued above, a correct (i.e., veridical) representation can be easily dismissed, since this construal would require that the luminance values of the stimulus could be analyzed to reveal the actual contributions of illumination, surface reflectance, and atmospheric transmittance that comprise the stimulus. Because this cannot be accomplished (cf. Section 1), there is no way to achieve a veridical representation of the scene. For similar reasons, a representation was also shown not to be assessable for correctness on a teleological approach, since the strong dichotomy between a representation and a misrepresentation required that *the* source could be represented.

Despite these arguments, it may still seem as though some notion of correctness is appropriate to ground the difference between a representation and a misrepresentation from a teleological perspective. One might therefore maintain that effects like those elicited by simultaneous brightness contrast are epiphenomenal in their character, an illusion or misrepresentation that occurs only in isolated circumstances and without any real behavioral efficacy. According to this view, in the vast majority of conditions the visual system represents the world correctly (in some sense of the term); only in a few instances—broadly cast under the net of "visual illusions" or "misrepresentations"—does the visual system fail.

The problem with this notion of "correct," however, is that it requires an explanation of our visual experience in the preponderance of cases. If it cannot be the actual conditions that generated the stimulus, then presumably it must be the properties of the stimulus itself—e.g., the luminance values projected on the retina. This might seem to be an adequate resting place, but it quickly generates problems of its own. One problem is that it now seems that what the visual system represents is the retinal stimulus—a position few, if anyone, would find attractive. A related problem is that it now makes the simultaneous brightness contrast effect even more perplexing. For since the visual system would now have to behave like a photometer by representing luminance, an explanation for *why* the system fails in the case of simultaneous brightness contrast must be offered.[18] And if the culprit has something to do with perceiving in terms of brightness rather than luminance, then this would seem to relegate all cases where brightness is perceived to the class of *misrepresentations*—in effect making *all* visual representations misrepresentations, since brightness perception is a fundamental aspect of vision. Unless one simply denies the simultaneous brightness contrast effect—a desperate move that merits

18. Specifically, an explanation must be offered for other, more problematic instances of brightness contrast that have different configurations from the one illustrated in Figure 3 (e.g., White's illusion, the Wertheimer-Benary stimulus, and the Inverted-T illusion, to name a few).

little attention—there seems to be no work remaining for a correct representation to do.

The crux of the matter, then, is that for a teleological framework like Dretske's to work it must be possible to delineate between correct and incorrect representations, since that is how functionally derived meanings explain the problem of representation. But since the inverse optics problem rules out this representational dichotomy, such a framework cannot offer either a naturalized account of visual content or an externalist justification for determining the content of visual states. In a very real sense, then, by moving the inverse problem to center stage (rather than ignoring or obscuring it) a different approach is required to explain visual representation. This is an important consequence, since it appears that an internalist position is also flawed in this regard (although for different reasons). The question then becomes, how can these two issues—naturalizing visual representation and content determination—be explained?

4. Teleosemantics Revisited. Dretske's teleological framework has been highlighted in part because it offered the most direct means for reaching the heart of the matter concerning visual (mis)representation. In this sense, it has been used primarily as a vehicle for establishing a negative conclusion. But this was not the only reason for focusing on Dretske's view; there is much in his argument that can be put to good use, despite the problems inherent in such an approach. The first step in preserving aspects of a teleological framework, then, would be to establish some of the more important merits of the position.

Perhaps the best reason has already been stated: using evolutionary history to define the representational parameters of the system (i.e., what it represents), avoids the intractable task of having to define visual representation in terms of the system's present state alone. Given the inverse optics problem, and the complexity of the task the visual system must contend with, appealing to evolved function is perhaps the only option. Evolved function by itself, however, cannot describe how the contents of visual perception are actually determined; for this, a perspective such as empirical ranking theory must be adopted (see Section 1).

For visually guided creatures, survival demands that complex spatiotemporal retinal stimuli are processed at rapid speed. That these patterns of light cannot uniquely determine the underlying structure of the world makes the representational capacities of the system remarkable, but not inexplicable. By generating visual representations according to how often a given retinal stimulus has occurred relative to all other instances of the same stimulus parameter in accumulated past experience, however, the efficiency by which the brain is able to represent the world *can* be understood.

On this view, the success achieved by the visual system arises by linking underdetermined retinal stimuli to their real-world sources empirically, on the basis of behavioral feedback accumulated over evolution and development. By instantiating such information over time in the circuitry of the system itself, vast amounts of trial and error would tend to accrue in the system, thereby providing the computational basis for representation and the determination of content in the context of the inverse problem.[19] As stated in Section 1, the evidence for this strategy comes from analyzing large natural scene databases, which serve to approximate the relationships between projected images and real-world sources that humans and other visual animals would have extracted behaviorally over time (Yang and Purves 2004; Howe and Purves 2005; Long, Yang, and Purves 2006, Wojtach et al. 2008). In adopting this tactic, however, a visual system could only come to generate *useful*—but not *correct*—representations. And although such representational states would be assessable, they are only assessable in terms of their biological usefulness to the organism.

If this is right, then the visual system can still be said to have evolved to represent the external world as in statement 1 above, much as Dretske argued for. But this fact alone does not offer any insight into the perceptual capabilities of the system as indicated in statements 2 and 3. The reason this was not an issue before was because one general notion of representation predicated on notions like natural signs and normal/abnormal conditions was deemed adequate for explaining visual representation and the determination of content. Thus, the real importance of the inverse problem is it demonstrates that if progress is to be made on the problem of visual representation and representational content, looking to what biological systems are supposed to represent (as expressed in 1 above) is not enough. In addition to this, the manner by which systems represent external conditions must be elucidated, since the contents of perceptual experience, as expressed in 2 and 3, will only be understood once this is done.

Consider once more the case of simultaneous brightness contrast pre-

19. While the details concerning the instantiation of such information in the system are still in their early stages, neurobiological support for the salience of the inverse problem comes from optical imaging of striate and extrastriate visual cortex, where the same patterns of neuronal activity have been shown to arise from retinal stimuli with different orientations, directions, and speeds (White, Basole, and Fitzpatrick 2001; White, Bosking, and Fitzpatrick 2001; Basole, White, and Fitzpatrick 2003). Rather than a "bottom-up" approach to vision, then, where patterns of neural activity are believed to be the combination of dedicated neural responses to orientation, direction, and speed, such evidence indicates that patterns of activity depend on more than what is present in a given stimulus. In terms of the argument offered here, this would be best understood in terms of the empirical relationships between images and sources that determined the evolution and development of the visual system.

sented in Figure 3. As previously stated, if the visual system could unravel the actual contributions that illumination, surface reflectance, and atmospheric transmittance had in generating the projected image, then the luminance values of the stimulus could be used to represent the world. But since this task cannot be accomplished—this just is the puzzle posed by the inverse problem—the visual system must be using a different strategy to create its representations. The approach of empirical ranking offers an explanation of how this might be achieved; it also points to a naturalistic explanation for the brightnesses (or any other subjective visual quality) that we experience on a routine basis.

As it pertains to Figure 3, then, we represent the scene in terms of brightness (not luminance) because the best way to produce biologically useful representations from underdetermined stimuli is with respect to the rank a particular stimulus has relative to the full range of conjoint luminance values that have occurred in the past, rather than by the actual luminance values on the retina at the present. The simultaneous brightness contrast effect arises, then, not because the system is failing to represent such circumstances correctly, but because an empirical strategy is relied upon to contend with the inverse problem. Although counterintuitive, this strategy can therefore explain why the contents of perception do not align with either the retinal stimulus or with the real-world conditions that generated the stimulus. In so doing, it provides a rationale for the discrepancies between appearance and reality that are most easily noticed in the case of the simultaneous brightness contrast effect and other so-called visual illusions.

This is also why it no longer makes sense to separate visual representations into two classes—those that are "correct" and those that are "incorrect" or "illusory"; instead *all* visual percepts are generated using the same empirical strategy (Purves and Lotto 2003; Wojtach 2005; Wojtach et al. 2008). It is not the case, then, that the visual system relies on this approach in some circumstances but not others; rather, the same framework is employed ubiquitously. The fact that this strategy cannot reveal reality (even though it may seem otherwise) only provides further support for the conclusion that the very ideas of correct representations and misrepresentations no longer apply.[20]

20. It might be thought that while a probabilistic approach can handle the problem of misrepresentation, it does not address the problem of "chronic" misrepresentation—e.g., perceiving a sunset as the Sun moving behind a motionless horizon. Rather than posing a problem, however, such chronic "misrepresentations" demonstrate the strengths of a probabilistic perspective in two ways. First, since representation is now to be cast in terms of biological utility (and not in terms of correct or incorrect), instances of chronic misrepresentation are better conceived of as instances where the biological utility is low. Indeed, such cases are not to be found in circumstances where

When conceived in these terms, it should be clear that a teleological perspective provides a way to describe *what* the visual system is representing—i.e., the sources in the world that generate visual stimuli. This is what the visual system was selected for; this is what its representations are supposed to be about. But such a perspective does not tell us much about *how* the system might actually accomplish this feat; moreover, it cannot predict—qualitatively or quantitatively—the perception elicited by a given stimulus, whereas a probabilistic framework like empirical ranking theory can (Purves and Lotto 2003; Wojtach 2005; Wojtach et al. 2008). If a theory of (visual) representation is measured by how well it can account for the relevant phenomena, then an approach that relies on past experience to link stimuli with sources by way of behavior enjoys much greater success.

Such considerations illustrate why a teleological perspective is inadequate, while offering a way to address how subjective experience can arise in a naturalistic framework. In particular, empirical ranking theory provides a way to recast 3 as

3*. what the visual system represents empirically is what we see.[21]

In addition, this manner of conceiving of visual representation also permits a way to recast statement 2. For now it can be stated that, because a biological system like the visual system

1. evolved to (is supposed to) represent external conditions,

the system can

2*. represent those conditions empirically, in terms of their biological utility.

---

the biological utility is typically high—e.g., the misperception of rapidly moving objects as slowly moving objects. Notice, too, that related instances of what might be considered chronic misrepresentation—e.g., the Sun appearing larger near the horizon than overhead (typically called the "Moon Illusion")—are better conceived in terms of a probabilistic model of visual space, since the apparent location of a source in three-dimensional space accords with the empirical rank of projected images arising from natural environments. Therefore, even these kinds of "misrepresentations" can be understood in terms of biological utility. Second, the problem of chronic misrepresentation assumes that the goal of vision is to reveal reality, whereas the present argument is that the goal of vision is to guide behavior by approximating reality according to the empirical rank of a retinal stimulus. In other words, if the goal of the visual system were to faithfully represent reality (including orbital mechanics), then presumably we would not perceive the Sun as setting; rather, we would perceive something more faithful to the rotation of the Earth and its relationship to the Sun.

21. Notice that what some might have viewed as a tautology in 3 is not present in 3*. Cf. note 17.

When visual representation is thought about in these terms, it is a stark reminder that a teleological framework and the notions of (proper) function and malfunction, correctness and error, are poor descriptions of what is actually taking place. Such notions should instead be replaced with a probabilistic framework like empirical ranking theory and the corresponding idea of biological usefulness, since doing so avoids the problems that are raised when conceiving of vision in terms of "correct" and "incorrect" representations. Notice, too, that just because this strategy lacks the notion of a "correct" representation, it would be wrong to assume that it lacks epistemic status. If a visual system is relying on the relative ranking of stimuli in past experience to represent the world, then this is an epistemic result if there ever was one. Just because it is a different epistemic result than expected (i.e., not one grounded on direct access to the world) is no argument against this model of representation.

5. Empirical Externalism. If progress is to be made on the issue of (naturalized) perceptual content, it should be clear that a teleological framework *can* ground a broad notion of representation. To explain the manner by which we perceive the world in the context of the inverse optics problem, however, more must be offered than just an appeal to teleological considerations. As might be expected, this shift in perspective has implications for an externalist position regarding content.

A standard reading of externalism maintains that the properties of a representational state are fixed by external factors. In the context of vision, this means that real-world sources are represented as having particular shapes, colors, and the like because those sources actually have such properties. What the inverse problem so vividly illustrates, however, is that this manner of thinking about visual representation cannot be correct. For while the relationship between real-world sources and their projected images are to be located externally, the underdetermination of retinal stimuli precludes any way of representing the true properties of those sources. The position offered here has argued that to contend with such circumstances, the visual system employs a probabilistic strategy based on the empirical rank of a stimulus to represent the world; but in so doing, the ability to represent the world correctly is traded for the ability to represent the world usefully. If this is accurate, then the content of a representational state cannot be fixed by a particular image–source relationship (cf. Figures 1–3); only accrued information from the past could do this. Therefore, because this position is significantly different from traditional views on representational content, yet it retains at least some of the virtues of a teleological perspective, let us call this view of representational content *empirical externalism*. This term indicates what is correct about externalism—the connection between sources and their pro-

https://doi.org/10.1086/597020 Published online by Cambridge University Press

jected images is based on relationships that occur external to the organism—while signaling that, in the context of the inverse optics problem, it is the *accumulation* of images over vast amount of trial-and-error experience that determines content. Without taking such factors into consideration, a naturalized explanation of the representational content of vision would be beyond our reach. This shift of perspective therefore provides a different way to conceive of the content of visual states, while avoiding the standard problems inherent in the internalism/externalism debate.

While this solution indicates how representational content can arise, a corollary of this argument is that the traditional problem of indeterminacy remains: although the content of representational states are fixed by accumulated experience, given the inverse optics problem there is simply no way to uniquely determine the content of the system's representational states as they pertain to the world. The best that can be achieved is the determination of content probabilistically in terms of a stimulus' empirical rank. This does not mean, however, that the way we represent the world is by chance: visual representations themselves are nonprobabilistic in their character, although they stem from a probabilistic system (cf. Wojtach 2005). What it does mean is that, at least with regard to vision, the problem of indeterminacy is a red herring, since there is no solution that *can* be offered. Rather than worrying about this result, however, it is time to simply accept it. For given the constraints of the inverse problem, as long as a visual system permits an organism to generate biologically useful representations of its environment by using past experience, then this is all that is needed to characterize the content of such states.[22]

Unlike other philosophical perspectives on representational content, a particular strength of the position offered here is that it is testable: if the empirical rank of projected images from natural scenes cannot predict visual experience, then this hypothesis about representation and representational content would be demonstrably false. In fact, however, empirical ranking theory can predict what observers perceive with a high degree of accuracy (Wojtach 2005), indicating the merit of this approach. Therefore, this argument supplies an extra dimension to the philosophical debate over representation by providing a specific, empirically motivated framework for how visual content can be realized.

Of equal importance, however, is the rationale for this stance: a probabilistic strategy like empirical ranking theory can explain the puzzling

22. A further issue pertains to whether the content in question here is best thought of in nonconceptual or conceptual terms. Although this issue falls outside the scope of the present argument, it is being currently explored in a separate manuscript on the topic of sensation and perception (in preparation).

features of visual experience—the so-called "visual illusions"—that other frameworks relegate to the status of mere anomalies. Indeed, while such perceptions seem to be maladaptive, adopting the change of perspective argued for here provides a way to consider such "anomalous" perceptions differently—namely as hallmarks (rather than limitations) of the visual system's ability to represent the world. In other words, the alleged misrepresentations of vision are better understood as the more obvious manifestations of an empirical strategy.

In sum, once the inverse optics problem is recognized and taken seriously, new challenges for theories of visual representation are uncovered, and old concerns can be dispelled. For when viewed from the perspective presented here, it becomes clear that assigning content to a visual representation cannot be achieved by an appeal to teleological function alone; nor can it be had by looking simply to the external properties of the world or the internal properties of the system. Rather, such content is fixed by the way the system has been shaped via past experience so as to generate biologically useful representations. So understood, the result is a new framework from which to consider the problem of perceptual content.

REFERENCES

Basole, Amit, Leonard E. White, and David Fitzpatrick (2003), "Mapping Multiple Features in the Population Response of Visual Cortex", *Nature* 423: 986–990.

Berkeley, George ([1709] 1975), "An Essay Towards a New Theory of Vision", in *Philosophical Works, Including the Works on Vision*. London: Everyman Press, 1–60.

Block, Ned (1986), "Advertisement for a Semantics for Psychology", *Midwest Studies in Philosophy* 10: 615–678.

Davidson, Donald ([1987] 2001), "Knowing One's Own Mind", in *Subjective, Intersubjective, Objective*. Oxford: Oxford University Press, 15–38.

Dretske, Fred (1981), *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

——— (1986), "Misrepresentation", in Radu J. Bogdan (ed.), *Belief: Form, Content, and Function*. Oxford: Oxford University Press, 17–36.

——— (1995), *Naturalizing the Mind*. Cambridge, MA: MIT Press.

Fodor, Jerry (1987), *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.

——— (1990), *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.

Gibson, James J. (1979), *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Helmholtz, Hermann von ([1866] 1924), *Helmholtz's Treatise on Physiological Optics*. Translation of the 3rd German edition. Edited by James P. C. Southall. Rochester, NY: Optical Society of America.

Howe, Catherine Q., R. Beau Lotto, and Dale Purves (2006), "Comparison of Bayesian and Empirical Ranking Approaches to Visual Perception", *Journal of Theoretical Biology* 241: 866–875.

Howe, Catherine Q., and Dale Purves (2005), *Perceiving Geometry: Geometrical Illusions Explained by Natural Scene Statistics*. New York: Springer.

Knill, David, and Whitman Richards (1996), *Perception as Bayesian Inference*. Cambridge: Cambridge University Press.

Laughlin, Simon B. (1981), "A Simple Coding Procedure Enhances a Neuron's Information Capacity", *Zeitschrift für Naturforschung* 36: 910–912.

Long, F., Zhiyong Yang, and Dale Purves (2006), "Spectral Statistics in Natural Scenes Predict Hue, Saturation, and Brightness", *Proceedings of the National Academy of Sciences USA* 103: 6013–6018.

Maloney, Laurence T. (2001), "Statistical Decision Theory and Biological Vision", in Dieter Heyer and Rainer Mausfeld (eds.), *Perception and the Physical World: Psychological and Philosophical Issues in Perception*. West Sussex: Wiley, 145–190.

Millikan, Ruth G. (1984), *Language, Thought, and Other Biological Categories*. Cambridge: MIT Press.

——— (1993), *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press.

Palmer, Stephen (1999), *Vision Science: From Photons to Phenomenology*. Cambridge, MA: MIT Press.

Purves, Dale, and R. Beau Lotto (2003), *Why We See What We Do: An Empirical Theory of Vision*. Sunderland, MA: Sinauer Associates.

Putnam, Hilary (1975), "The Meaning of 'Meaning'", in *Mind, Language and Reality: Philosophical Papers,* vol. 2. Cambridge: Cambridge University Press, 215–271.

Rao, Rajesh P. N., Bruno A. Olshausen, and Michael S. Lewicki (eds.) (2002), *Probabilistic Models of the Brain: Perception and Neural Function*. Cambridge, MA: MIT Press.

Simoncelli, Eero P., and Bruno A. Olshausen (2001), "Natural Image Statistics and Neural Representation", *Annual Review of Neuroscience* 24: 1193–1216.

Stocker, Alan A., and Eero P. Simoncelli (2006), "Noise Characteristics and Prior Expectations in Human Visual Speed Perception", *Nature Neuroscience* 9: 578–585.

Weiss, Yair, Eero P. Simoncelli, and Edward H. Adelson (2002), "Motion Illusions as Optimal Percepts", *Nature Neuroscience* 5: 598–604.

White, Leonard E., Amit Basole, and David Fitzpatrick (2001), "Effect of Speed on Responses of V1 Neurons to Motion of Contour and Terminator Cues", *Society for Neuroscience* Abstract 27, Program 164.4.

White, Leonard E., W. H. Bosking, and David Fitzpatrick (2001), "Consistent Mapping of Orientation Preference across Irregular Functional Domains in Ferret Visual Cortex", *Visual Neuroscience* 18: 65–76.

Wojtach, William T. (2005), *An Empirically Constrained Study of Visual Perception*. PhD Dissertation. Durham, NC: Duke University.

Wojtach, William T., Kyongje Sung, Sandra Truong, and Dale Purves (2008), "An Empirical Explanation of the Flash-Lag Effect", *Proceedings of the National Academy of Sciences USA* 105: 16338–16343.

Yang, Zhiyong, and Dale Purves (2004), "The Statistical Structure of Natural Light Patterns Determines Perceived Light Intensity", *Proceedings of the National Academy of Sciences USA* 101: 8745–8750.