Journal of Philosophy, Inc.

Biological Functions and Perceptual Content

Author(s): Mohan Matthen

Source: *The Journal of Philosophy*, Vol. 85, No. 1 (Jan., 1988), pp. 5-27

Published by: Journal of Philosophy, Inc.

Stable URL: <http://www.jstor.org/stable/2026898>

Accessed: 28/05/2009 15:12

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=jphil>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.

Journal of Philosophy, Inc. is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*.

<http://www.jstor.org>

THE JOURNAL OF PHILOSOPHY

VOLUME LXXXV, NO. 1, JANUARY 1988

BIOLOGICAL FUNCTIONS AND PERCEPTUAL CONTENT*

THE computational approach to perception has become increasingly prominent over the last decade or so. The recent publication of David Marr's *Vision*¹ makes computational theories accessible to philosophers, who may now judge for themselves how unified, wide-ranging, powerful, and intellectually satisfying the program has proved to be.

A striking feature of computational theories of perception, from the philosophical point of view, is that they employ the intentional characterizations of perceptual states insisted upon by Edmund Husserl and other phenomenologists. Most other scientific treatments of perception concentrate on the external causes of perception and how they result in the stimulation of the sensory organs, the neurological organization of perceptual systems, and the discriminatory resources afforded to organisms by the perceptual systems. Though not ignoring these important topics, the computationalists characteristically regard perceptual states as representing the external world and seek to explain them under this rubric. Marr claims, for example, that vision represents three-dimensional surfaces as composed of "generalized cones," and tries to show how such representations may be derived from representations of two-dimensional outlines of the sort found in line drawings and how the latter may, in turn, be extracted from a pattern of optical stimulation present on the retina.

Most philosophers today, and most computational theorists, are materialistically inclined: they are not inclined to allow that there

* A version of this paper was read to the philosophers at Dalhousie University. I am grateful to them all for lively discussion, and to Terry Tomkow and Bob Martin in particular for helpful suggestions. I owe no small debt to the perception reading group in my department, and especially to Bruce Hunter.

¹ New York: Freeman, 1982.

could be a causally effective aspect of a perceptual state that is not material.² For this reason, any widely acceptable foundational account of computational theories of perception must explicate what it means to endow a material state with representational content. It is my purpose in this article to develop a theory of perceptual content which meets this requirement.

An important constraint on such a theory results from the fact that computational theories tend to attribute *distal* content to perceptual states. They claim that we perceive things like rigid bodies, motions in three-dimensional space, and reflectances of surfaces: these are entities in the external world. Thus, the theory we are looking for cannot identify content with some purely internal concomitant of a perceptual state—it is important, in other words, to avoid the materialistic analogue of sense-datum theories, which claim that the content of a state is properly specified by the introspectible feel that accompanies it.

It is often assumed that the alternative to “internalist” theories of perceptual content is a causal theory. I shall argue in section III that this too is a false lead. Sections II and III prepare us for a new approach to perceptual content. The theory of perceptual content I shall propose is based on biological functions. It employs a notion of *biological adaptation* which results from an extension of Larry Wright’s work on functions (sections IV and V). I shall show (section VI) how chains of adaptations are useful in accounting for the attribution of content, first, to states that may be regarded as the evolutionary precursors of perception and, then, to the full-blown perceptual states that humans experience (section VII). I shall argue that the content-attribution scheme I propose accounts for distality of content and, indeed, that it can form the basis of *realism* with regard to perception (section VIII). Finally, I shall make a suggestion how the content-attribution scheme might be able to account for the point and utility of intentional attribution in general (section IX).

I. WHAT IS PERCEPTUAL CONTENT?

According to Husserl, a perceptual state is *about* something, its *object of perception*. It achieves this aboutness by means of its *noema*. According to one influential interpretation, a *noema* stands to the object of perception much as the Fregean sense of a referring expression stands to its reference.³ This conceptualization of the con-

² Qualia seem to constitute an exception: many psychologists talk as if they exist and influence our actions.

³ The interpretation is due to Dagfinn Føllesdal: see his “Husserl’s Notion of Noema,” this JOURNAL, LXVI, 20 (Oct. 16, 1969): 680–689. The classic source of Husserl’s own exposition is his *Ideas* (1913), W. R. Boyce Gibson, trans. (London: Macmillan, 1931).

tent of a perceptual state is, I believe, particularly appropriate to explaining the structure of computational theories. Let us say that the content of a perceptual state, its Husserlian *noema*, is *expressible* by means of a definite description. Then the object of perception is the unique object that satisfies the definite description. If there is no such object, the perception is nonveridical—it is a misperception.⁴

The content of a perceptual state may be regarded as having two components. First, it *presents* an object as possessing a certain characteristic—it may present an object *as* red, for example, or *as* moving in a certain direction, or *as* five feet away. Second, there is something about the state which accounts for the fact that the state has a unique object—not just any red thing is the object of a red-presenting state, but some unique red thing. An adequate theory of perceptual aboutness must account for both these features of perception, its *mode of presentation* and its *individual-directedness*.

In my opinion, any adequate theory of the individual-directedness of perception will be causal. A promising theory along these lines has recently been offered by David Woodruff Smith, elaborating on an idea of John Searle's.⁵ According to this theory, the definite description that expresses the content of a perceptual state has a self-referential causal-indexical component. Searle's idea is that the content of a perceptual state is always expressible in this form: *the individual that has the property F and is causing THIS state*. The object of a perceptual state, then, will be whatever satisfies this description.

The only variable element in content so expressed is the property *F* by which the individual is characterized. Every perceptual state is directed toward some individual causing the state itself, but some states are directed toward the *red* object causing it, some toward the *moving* object causing it, and so on. This variable element is the *mode of presentation*. Computational theories are efforts to explain how perceptual states present external objects. The question I shall seek to address in what follows is this: What feature of a perceptual state makes it a presentation of a particular property, say redness? As a warning against a possible confusion, I should remark that, even though content as a whole has a causal-indexical element, it does not

⁴ This neglects what David Lewis has called "veridical hallucinations": see his "Veridical Hallucination and Prosthetic Vision," *Australian Journal of Philosophy*, LVIII (1980): 239–249. This is a problem that concerns the formulation of what I shall call the *individual-directedness* component of the *noema*, not the *mode of presentation*, and so it will not concern us in this paper.

⁵ Searle, *Intentionality* (New York: Cambridge, 1983), pp. 61–71; and Smith, "The Ins and Outs of Perception," *Philosophical Studies*, XLIX, 2 (March 1980): 187–217.

follow that an account of the mode of presentation should be causal as well.

II. INTERNALIST CRITERIA OF CONTENT

Some epistemologists have claimed that the quale, or subjective character of a state, constitutes its mode of presentation—thus a state presents an object as red in virtue of its being associated with a red quale or sense datum. Others give some internal material state the same role as qualia. Just as pain has been identified with a stimulation of the C-fibers, a red-presentation might be identified with some neurological concomitant of seeing red. Such theories are *internalist* because they ascribe content to states solely on the strength of their intrinsic character and refer to nothing outside the perceiver, especially not to the external things represented by perception.

An internalist theory is unsatisfactory for several reasons. Most importantly for our present purposes, it cannot adequately account for the distality of perceptual content found in computational theories. As an example of what I mean, let us consider in some detail Edwin H. Land's theory of color vision,⁶ which is regarded by many as an early paradigm of the computational approach to perception.

The cone cells of the retina are sensitive to the color of light. Consequently, many psychologists and many philosophers have assumed (in a tradition going back to Newton's *Optics*) that the color we see is the color of the light incident on our retinas. But this is false. Our perceptions of color are, in fact, much more closely correlated with the true color of surfaces, even when they reflect light that might be expected to be quite misleading. This was recognized to be the case as long ago as Helmholtz, but most conclusively established in a series of experiments, conducted by Land and J. J. McCann, in which human observers were shown to be able to detect that two surfaces were different in color, and what their true colors were, even though the illumination was so rigged that the surfaces were sending light of exactly the same color to the eye and so should have been indistinguishable.

But how can vision detect true color if it is sensitive only to light color? The color of light reflected from a surface is a function of two independent variables—the color of the surface itself, and the color of the illumination incident upon it.⁷ Thus, the color of light incident

⁶ "The Retinex Theory of Color Vision," *Scientific American*, CXXXVII, 6 (December 1977): 108–128; and Land and J. J. McCann, "Lightness and Retinex Theory," *Journal of the Optical Society of America*, LXI (1971): 1–11. Marr (17) says that Land was the first to pose a "clear computational question" about perception.

⁷ Let us divide visible light into three wavelength ranges: long, medium, and short. Then the color of a surface can be defined as the ordered triple of the

upon the retina is compatible with an infinite number of covarying combinations of surface color and illumination color. What is needed, then, from an adequate theory of color vision, is an account of how the visual system is able to discriminate surface colors, given only light color; in other words, how it separates out the contributions to afferent light color due to each external factor.

In the world in which we live, there are many abrupt changes in surface color, but very few abrupt changes of illumination. According to Land, color vision uses a process that is effective because of this fact. It distinguishes types of brightness gradient in the image incident upon the retina. Sudden changes of brightness are identified as being due to changes in surface color, and more gradual changes are identified as being due to changes in illumination. In this way, our vision assigns a surface-color *gradient* to each point in a scene. *Absolute* colors are then assigned to points by a process of calibration, the details of which need not concern us here.

It is very important to be clear on what Land uses to explain what. First, Land uses psychological experiments to identify a physical variable, surface color, as that which is presented by color vision. The problem he then tackles is that information concerning this detected property is unavailable at the retina. So Land appeals to an *indicator* of the detected property which he claims *is* present in the retinal image, and uses it to explain how the visual system has access to the presence of the external property.

Obviously this procedure raises a number of interesting questions. What is important for our immediate purposes is that Land type-identifies perceptual states by the external property they present, not by their quale nor by any other internal characteristic. The direction of explanation in Land's theory is from external properties and their indicators to the content of states. Internalist theories reverse this direction of explanation. For them, the similarities relevant to content attribution are internal, not external, and any representational relationship borne by states to external reality must be accounted for by internally specified manipulations of content. As an extreme of this attitude, consider qualia theories. They type-identify states by their qualia, and identify external objects as having the same property because they tend to produce qualitatively similar states in us.⁸

reflectances of the surface for light in these three wavebands. The color of *a beam of light* is the average of wavelengths in the beam weighted by the energy present in each wavelength.

⁸ There are two versions of a definition of color by means of qualia. The first sort makes color a construction out of qualia. The second sort uses qualia to identify a

Now, it is possible to argue that Land's type-identification will have to be nearly *coextensive* with a type-identification by internal criteria. For, as we have seen, the experimental verification of Land's hypothesized mode of presentation lay in the ability of observers to identify the true colors of surfaces (even when they sent misleading signals to the eye). One would think that observers identify a surface as having a color as a result of *some* internally available feature of the perception, perhaps a *quale*. So there must be *some* intrinsic difference between perceptions of different things. To put this point in general terms, a perceiver's *epistemic access* to the content of his or her perceptual states must be through internal characteristics. This access has to be reliable, if not infallible; so there will always be an extensionally correct (or nearly correct) internalist definition of content.⁹ Nevertheless, such definitions will misrepresent computational theories, because they will miss what such theories take to be the essence of content.

III. CAUSAL THEORIES OF THE MODE OF PRESENTATION

There are two sorts of causal theory. The first kind, which I shall call *outside-in*, attributes content to a perceptual state on the basis of what causes it. Broadly speaking, a theory of this type will allow that two states with the same intrinsic (i.e., internally specifiable) character may have different content, because two states with the same character may have different causal origins.

The shortcomings of the outside-in approach have been decisively exposed by Jerry Fodor.¹⁰ If a visual state counts as being of a rigid body simply because it was caused by one, then the content of the state will be unhelpful in explaining the use a perceiver makes of it. A visual state might have been caused in me by a rigid body without that fact being known to me. My *treating something as* a rigid body must, then, be explained by invoking an additional state, my *perceiving it as* a rigid body. But, if the latter state is invoked, the fact that the

real property. [An example of the second sort: red is the external property that causes this sort of quale in me. For an ingenious discussion of this sort of definition, see Edward Wilson Averill, "Color and the Anthropocentric Problem," this JOURNAL, LXXXII 6 (June 1985): 281-304.] It is the first sort of definition which is criticized in this section. The second sort is criticized, in effect, in the discussion of Keith Campbell and C. L. Hardin in section VIII below.

⁹ Some philosophers would contest the value of such a definition, however. It does not follow from my argument that there will be *inter-personally applicable* internalist definition of content. For qualia theories, the classic argument that purports to show such a definition false is the inverted-spectrum argument. There is an obvious materialistic analogue of this argument: presumably, token-token identity theorists would endorse this analogue.

¹⁰ "Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology," *Behavioural and Brain Sciences*, III, 1 (1980): 63-73.

former is or is not caused by a rigid body will be otiose—what I do will be the same in either case.

An *inside-out* theory of perceptual content recognizes that there must be internal features of a state which provide the perceiver with access to its content. It thus looks at intrinsic features of a perceptual state and asks: How does a state with these intrinsic features present its object? In this way, it assigns distal content to *internally identifiable types* of states.

Now, suppose that we have two states with the same intrinsic character, one caused by a particular external object or feature, and the other not. An inside-out theory will assign the same content to both states, and so this assignment must not imply that the external object or feature is present. There are two options available for achieving an assignment of this sort.

First, one may retain the assumption that content is cause, and so retreat to some more proximal common cause of the two states. For example, when it was believed that two different surface colors could result in the same sort of state, under different conditions of illumination, philosophers claimed that it is not surface color that we see, but rather the color of the light incident upon the eye. Similarly, it has been held that what we see is projected shape, size, and motion, not actual shape, size, and motion—the move is familiar. It is this move which we wish to avoid in order to maintain the characteristic distality of content attribution in computational theories.

The other alternative is to drop the causal assumption. This approach need not be as radically at odds with the general causal intuition as it looks, for one might say that a state with content *C* is *usually* or *normally* caused by an object that is *C*. My argument here seeks to show that no such move will work. Causal theories of the mode of presentation have to be abandoned altogether.

Let us look at the matter in more detail. We have seen that an inside-out theory of perceptual content attempts to assign distal content to internally identifiable types of states. If a causal theory is seen as the way to account for distality, an inside-out theorist will have to define content in terms of what causes a state of a certain sort. There are two versions of such a theory, depending on whether we define content by the *invariable* or the *normal* causes of a kind of state:

IOC: A perceptual state with (internal) feature *S* will have mode of presentation *F* if *S*-states are invariably/normally caused by things that have the property *F*.

It is a familiar fact that theories of content which appeal to the invariable causes of kinds of states cannot deal with misperception.

For, if a state that presents a thing as *F* is always caused by a thing that *is F*, then it cannot come into being unless it is veridical.¹¹ I shall argue that theories that appeal to the *normal* causes of misperception will not work either: for they cannot deal with what I shall call *normal misperception*.

Before we find out what normal misperception is, let us briefly look at two sorts of misperception which *are* adequately accounted for by the version of IOC which appeals to the normal causes of a sort of perceptual state. The first is caused by *malfunction*: if a system is not working normally, because of genetic damage or injury, then misperceptions may result. Second, *maladaptation*: if a system is exposed to environmental circumstances to which it is not adapted, circumstances manufactured by a laboratory worker for example, then it may be "fooled." Is misperception always due to malfunction or maladaptation, as Fodor conjectured in "Semantics, Wisconsin Style"? If so, then, on a plausible understanding of the 'normally' in IOC, we should have an adequate causal account of the mode of presentation.

But there is another source of misperception. Perceptual states play a mediating role between external circumstances and appropriate behavior. But the features of the external world which affect a system may not be those which are "task-relevant." This is immediately apparent in the case, already described, of color vision. Our retinas are sensitive to the color of light. But light color tells us little that is useful about the things we have to deal with; so we are in need of a system that is able to process afferent "sensory information," yielding a characterization of the external world in terms of surface color.

Now, in carrying out this "processing," the visual system is hampered by the fact that the afferent sensory information it works with *underdetermines* conclusions about surface color. That is, the afferent information is compatible (given the laws of how light is reflected by surfaces) with several different and incompatible "conclusions" about surface color. So color vision uses an indicator of surface color which is present in the retinal image to arrive at the representation the organism needs. However—and this is the important point—the indicator is not perfect. It is possible to create a surface on which the color changes gradually, in just the way that illumination normally changes. In uniform illumination, such a surface would be perceived as uniformly colored and variably illuminated.¹² Here we have a

¹¹ See Fodor, "Semantics, Wisconsin Style," *Synthese*, LIX, 3 (June 1984): 239–250.

¹² Imagine a white table-cloth illuminated from one end by a white lamp and from the other by a red lamp. The illumination will vary gradually across its expanse,

situation in which the visual system represents a surface-color gradient, nonveridically, as an illumination gradient; in other words, we have a misperception.

The kind of misperception I want to draw attention to arises from the fact that our ability perceptually to discriminate the colors of surfaces in widely varying conditions of illumination is made possible by the use of indicators that are imperfect, but nevertheless the best available. If the visual system did not use such indicators, the only kind of color we perceived would be the color of light incident upon the retina, which is useless. The price of using the imperfect correlation is occasional misperception, a small price to pay for a useful representation of the world.

It is obvious that the misperceptions arising from the use of such imperfect indicators will not be traceable to *malfunction*: the error lies in the indicator itself, not in how it is used. Nor can the misperception be claimed to result from the visual system's being *mala-dapted* to the situations it encounters; for the indicator is not perfectly reliable even in the range of situations to which the visual system is adapted. (In many such cases, no indicator could be, since the afferent information underdetermines the detected property.) The point that has to be understood is that, even with reference to the range of situations to which the system is well adapted, it is better to use an imperfect indicator than to have no access to task-relevant properties. To use veridicality as a filter on this range of situations is to show a touching, but quite unbiological, devotion to truth.

It is misperceptions of this sort which I shall entitle *normal*. They result from the use of indicators that are imperfect but the best available, where this use confers an over-all advantage upon the organism, despite the occasional occurrence of error. Such misperceptions cannot easily be accommodated by a causal theory, because these are cases where, for example, a perception of red is caused by a blue surface even in normal circumstances.¹³

going from nearly white at one end to red at the other. If Land is right, this cloth should not look white at one end, red at the other. It should look uniformly white, and red-lit at one end, and white-lit at the other. Now, take a color photograph of the cloth and print the photograph on the cloth. The cloth will now really be red at one end, shading gradually to white at the other. Now, replace the red lamp by a white lamp of the same intensity. The cloth will look uniformly white, and as before, variably illuminated.

¹³ Normal error is the subject also of Matthen and Edwin Levy's "Teleology, Error, and the Human Immune System," this JOURNAL, LXXXI, 7 (July 1984): 351-372. The concept also emerges clearly in John Haugeland, "Semantic Engines: An Introduction to Mind Design," in his *Mind Design* (Cambridge, MA: MIT Press, 1982), pp. 1-34 (see especially 17/8); and in Christopher Cherniak, "Computational Complexity and the Universal Acceptance of Logic," this JOURNAL, LXXXI, 12 (December 1984): 739-758.

IV. BIOLOGICAL FUNCTIONS: WRIGHT'S ACCOUNT

The first step toward outlining my theory is to introduce the notion of a *biological adaptation*. Functional explanation in biology seeks to explain why an organism has a particular characteristic (or disposition) by appeal to other characteristics it has as a consequence. There is a good example in the "Science and the Citizen" section of the January 1982 issue of *Scientific American*. There, under the heading "Thin Red Line," we find a report of speculation concerning the function of the adherence of red cells to immune complexes in mammals. According to this report, the adherence phenomenon had been well known, but thought to have no "functional significance." But it had recently been discovered to be modulated by time and circumstance; it is depressed in human cancer patients, and enhanced in humans with autoimmune diseases. These observations led to two hypotheses, due to different teams of investigators: one team claims that "red cells must be the main agents for clearing immune complexes from the circulation," and the other that they "serve as intermediaries that promote the interaction of a T-cell and its specific antigen, thereby enhancing T-cell activity." The question asked here is of the form: Why does the immune system have complexes that do *F*? And the answer takes the form: Because as a result of having such entities, the system is able to do *G*.

Larry Wright provides us with a good foundation for understanding why explanation by reference to functions is applicable to biological entities.¹⁴ His analysis of the meaning of 'Z is the function of X' runs as follows:

- (a) X is there because it does Z, and
- (b) Z is a consequence (or result) of X being there (161).

In the case of biological entities, Wright tells us to evaluate clause (a) above by reference to natural selection. Roughly the idea is this: if the consequence Z is what was responsible for X's evolving or is what is responsible for X's persisting through continued evolutionary change, then Z is the function of X. Otherwise it is not.

Wright's account is distinguished from others by its introduction of a *historical* element into the analysis of function ascription. It is this historical element that enables him to show how a consequence of X's presence can explain X's presence. Consider, for the sake of contrast, a theory that lacks the historical element. Robert Cum-

¹⁴ "Functions," *Philosophical Review*, LXXXII, 2 (April 1973): 139–168; cf. also John Bigelow and Robert Pargetter, "Functions," this JOURNAL, LXXXIV, 4 (April 1987): 181–196.

mins¹⁵ tells us that functional explanation regards a function as a complex disposition that is explained by showing how simpler dispositions mesh together to result in the complex disposition. Cummins's theory gives an admirable account of the "forward" explanatory direction from *X* to its function, but not of the "backward" direction from the function to the characteristics by which it is achieved. Consequently, Cummins counterintuitively denies that there is any explanation in the backward direction. But it is precisely this backward explanatory direction which is accounted for by the historical element in Wright's analysis.

The historical element in Wright's account is important also because, if Cummins is right about functions, then there will be no way to choose between rival function ascriptions. We saw earlier that there were two rival hypotheses concerning the "functional significance" of red blood cells' adhering to immune complexes—that it contributes to the removal of immunological detritus, and that it enhances T-cell activity. Now suppose, as is possible, that in fact the adhesion phenomenon does both these things. Then Cummins is unable to choose between the rival hypotheses, because both show how the phenomenon contributes to some disposition of the immune system. Yet, this is a common phenomenon: many features of an organ are such as to have multiple effects, only some of which are functionally significant. What is to prevent us from giving a Cummins-style functional analysis of all or any of these effects?

This consequence of Cummins's view is particularly counterintuitive when it involves effects that actually reduce fitness. The narwhal has a large tusk which actually reduces its mobility. Presumably the tusk is there in order to increase the conspicuousness of a narwhal in mating situations. Obviously, then, the function of the tusk is to increase the sexual attractiveness of the animal; certainly its function cannot be to reduce mobility. Yet, a Cummins-style functional analysis can be given of how the tusk reduces the narwhal's mobility. Wright's analysis is superior to Cummins's because it is able to discriminate between effects on the basis of which one was historically responsible for the phenomenon that is to be functionally explained.

V. BIOLOGICAL ADAPTATIONS

Although I view Wright's analysis as basically correct, it needs to be extended. It is not only the result of an activity that plays a historical role in selection, but also the surrounding circumstances that make the selected activity a good way, or the only way, to achieve the result.

¹⁵ "Functional Analysis," this JOURNAL, LXXII, 22 (Nov. 20, 1975): 741–765.

Consequently, functional explanations found in biology make essential use of facts other than those mentioned in Wright's account.

A good example is to be found in a recent account of ribosome synthesis in the bacterium *E. coli*.¹⁶ According to this account, the bacterium needs to detect and control the rates at which it synthesizes ribosomes, and the proteins out of which ribosomes are built up. The author reasons as follows: any bacterium acts so as to maximize cell growth and cell division. Now, in order to maximize cell growth and division, the cell has to make enzymes. But enzymes can be manufactured only by ribosomes, so the bacterium must manufacture more ribosomes. On the other hand, the ribosome and its components are energetically expensive to manufacture, and making a surplus is wasteful. Therefore, the bacterium must detect and control the rate of ribosome synthesis and ribosomal protein synthesis, making sure that it is adequate but not excessive.

This exposition endows the activities of the organism not only with functions, but with functional structure. It makes implicit use of functional statements such as:

The function of enzyme synthesis is to facilitate increased metabolism and reproduction.

The function of ribosome synthesis is to facilitate enzyme synthesis.

But it also alludes to certain facts that make some activity *necessary* for the optimal performance of some function. These facts may concern the internal organization of *E. coli* or its parts:

Ribosomes are sites for enzyme synthesis.

Or they may concern *E. coli's* environment (in a broad sense of that term):

Ribosomes are energetically expensive to manufacture.

The role these facts play in the explanation is to subordinate one function to another. The function of ribosome synthesis is subordinated to the function of enzyme synthesis, because the latter is necessary for the former *given that ribosomes are the sites at which enzymes are manufactured*. The function of controlling the rate of ribosome synthesis is subordinated to the *proper* performance of ribosome synthesis, because ribosome synthesis is expensive. By 'proper', I mean "such as not to detract from the performance of other functions." Clearly, some measure is required by which to

¹⁶ Masayasu Namura, "The Control of Ribosome Synthesis," *Scientific American*, CCL, 1 (January 1984): 102-114.

balance off one performance against another: this measure is fitness. In short, facts are used to subordinate one function to another as means to end.

How are we to account for the explanatory value of these facts? The most natural way is to extend Wright's analysis of functions so as to allow that certain facts might have been historically influential in shaping the present constitution of an individual. Thus, consider two activities *F* and *G* of an organism such that

- (1) *G* has the consequence *F* in circumstances *C*.
- (2) *G* is there because it does *F* in circumstances *C*.

Under these conditions, we shall say that *G* has function *F* in circumstances *C*; or, conversely, that *G* is the realization of *F* in *C*. The distinction between activities (*G*), the functions they serve (*F*), and the circumstances in which *G* is subordinate to *F*, will turn out to be crucial to our account of perceptual presentation.

VI. THE ORIGINS OF PERCEPTION

The above account of functions is adequate for understanding claims about standing features of organisms, e.g., "The function of ribosome synthesis is to facilitate increased metabolism and reproduction." But there is an application of functional explanation which we have yet to touch upon. To see why, consider that biological entities such as fingernails and eyebrows have functions. But fingernails do not, in order to perform their functions, modify their activities to suit changing circumstances—they just sit there.

By contrast, *E. coli*'s control of its ribosome synthesis is a much more complex activity: it involves responding differently to different situations. When ribosomes are being manufactured faster than they can be used, there will be idle ribosomes present. These attach to a gene (an *operon*), which then shuts down further ribosome synthesis. When all available ribosomes are involved in enzyme manufacturing activity, the operon becomes unattached again, and ribosome synthesis starts up again. A similar mechanism controls the manufacture of the ribosomal components. Thus, the bacterium's behavior is *plastic*.

It is obvious that the plastic behavior of *E. coli* is subordinate to its functions in exactly the same way as some of its functions are subordinate to others. The control *system* is a permanent feature of the bacterium, and its presence is explained by the need to make neither too many nor too few ribosomes. Given this explanation of the existence of the control system, it seems possible to extend the chain of means and ends down into the plastic activities involved in ribosome synthesis. Let us distinguish between two levels of description. A system can be described either as executing a constant function or as

behaving in a variable manner. These two levels are linked by the fact that, in different sets of circumstances, the system has to perform different activities in order to be executing the constant function. Thus, we have the familiar pattern of variable activity being subordinated to a function, and being variable simply because the facts that mediate the subordination are themselves variable.¹⁷

Now, it is obvious that Wright's historical stratagem will be of no use in explaining this pattern of subordination, since functionally subordinate activities and subordinating circumstances covary concurrently. What is required here is not historical causation, but direct causation. There are two patterns of direct causation which could be used. The system could be taken to be functionally so organized that changing circumstances directly influence functionally directed behavior. Or it could be taken to be functionally so organized as to have internal states which *register* changing circumstances and which then initiate the activity functionally appropriate to the circumstances thus registered. Broadly speaking, the latter will be used when there is a possibility that behavior appropriate to circumstance *C* will be initiated even when *C* does not obtain—the resulting inappropriate behavior may then be explained by the system's *falsely* registering *C*. I shall suppose that such internal states are the evolutionary precursors of perceptual states, and ask what would be a suitable schema for attributing content to them. I do not mean to imply that such content attribution is legitimate. My interest in pursuing a content-attribution schema for such states lies solely in the extension of the schema to perception.¹⁸

We have seen that our bacterium *E. coli* does indeed have an internal state that mediates functionally appropriate activity in changing circumstances. When there are idle ribosomes they attach to an operon. When in this state, the operon "switches off" ribosome synthesis. When the ribosome is needed, it gets detached from

¹⁷ Some philosophers have treated variable activity, under the rubric of "goal-directed activity," quite separately from functions—the tendency is to give an account of these in terms of concurrent causation only. See Matthen and Levy (*op. cit.*, note 13), section II, for a criticism of such theories. It is a virtue of the current account that it eludes these criticisms, which are also dependent on the notion of normal error.

¹⁸ In his "Intention and Action among the Macromolecules," in N. Rescher, ed., *Current Issues in Teleology* (Lanham, MD: University Press of America, 1986), pp. 65–76, Alexander Rosenberg shows that biochemists have a tendency to describe such states in intentional terms, quoting from a biochemistry textbook which, in the space of a single paragraph, describes tRNA as recognizing, discriminating, undertaking tasks, making errors, editing, correcting, avoiding, and proofreading (65/6). Rosenberg argues that the analogy between goal-directed behavior "among the macromolecules" and perception, intention, and action in humans is enough to justify or even compel the use of intentionality. I do not agree with this.

the operon. In this new state, the operon switches on ribosome synthesis. Let us entitle these states of the operon *quasi-perceptual states*. The question is: What is their content? Intuitively, the right answer is given by the fact that the functional role of the states is to initiate the activity appropriate to the circumstances. The control system has to "sense" when ribosome synthesis is going too fast or too slow, and these are the states which do it. So the ribosome-attached state of the operon is a sensing of excess ribosomes, and the ribosome-free state a sensing of the need for more ribosomes. That is the only answer that links the subordinate behavior to its function in the right way.

This is just where the causal theory has found a foothold; for these states are *caused*, respectively, by excess (i.e., idle) ribosomes and by there being no excess ribosomes. But we have seen that there is reason to reject the causal account, as applied to real perceptual states, because it is incapable of accommodating normal error. Another difficulty for the causal theory is that causes can be re-described. The states in question can be said to be caused not just by an excess or deficiency of ribosomes, but by the presence or absence of *idle* ribosomes or of molecules with a certain chemical composition, and so on. The causal route does not, in other words, give us a schema that is potentially able to handle the intensionality of perceptual content.

There is another way of getting to what we have identified as the "right answer." Why not use the functional hierarchy to identify the mode of presentation? Might it not be because the *function* of the state is to trigger the action appropriate to an excess of ribosomes, that we want to say that its content is precisely the presence of such an excess? This way of identifying content is going to yield results quite close to the causal theory, simply because one should expect that, if a state has the function of triggering the appropriate response to the presence of Xs, then it will be brought into being by Xs—otherwise, there would be no reliable matching of response and situation, and natural selection would have to "search for" some other way of achieving the appropriate response. In fact, the two theories will diverge precisely in those cases where there is misperception. For, if a state that has the function of bringing about the appropriate response to the presence of xs does come about in the absence of xs, the functional theory would lead us to say that the organism falsely sensed the presence of xs. But this is precisely what we ought to be led to say, is it not? After all, we have given these internal states a role in the theory precisely in order to explain the system's responding to the presence of xs, even when xs are absent.

These reflections suggest the following schema for ascribing a "mode of presentation" to the quasi-perceptual states we are considering:

Q: A quasi-perceptual state is a state that has the function of bringing about the functionally appropriate response to the presence of things of a certain sort, where things of that sort are sometimes present, sometimes absent, and

A quasi-perceptual state with internal characteristic *S* is a presentation of property *F* if and only if the function of *S*-states is to bring about the functionally appropriate response to the presence of things that are *F*.

Note that Q is an inside-out schema for identifying the mode of presentation and that, although it makes the assignment of content to a *type* of state presuppose an external world, it does not demand that the assignment of content to an *individual* state imply the existence of external objects. In this way, *Q* partakes of a weak form of what Putnam has called "methodological solipsism."

VII. TOWARD PERCEPTION

Schema Q has one very obvious disadvantage when it is applied to real perception such as that which we humans experience. Quite simply, there is no such thing as *the* proper response, or even a range of functionally appropriate responses, to what perception tells us. Suppose I see my glass empty. I may not even come to believe that my glass *is* empty, since it is possible that my perception is misleading. And, even if I do come to believe that my glass is empty, what I shall do in response to this depends very complexly on my principles, my weaknesses, my desires, and so on, and on my beliefs concerning the state of the world. What time is it? How many have I already had? Obviously, then, schema Q is not going to work: there is no functionally appropriate response mediated by perception.

The most obvious way to retain the intuitions behind the functional schema is to get rid of the reference to the proper behavioral output of perception. Suppose we think of perception not as leading directly to action, but as leading to epistemic and practical assessment, which, in turn, leads to action. Can we not then define the output of perception in terms of its input to these faculties? The following is an attempt to do this:

P: A perceptual state is a state that has the function of *detecting* the presence of things of a certain type, where things of that type are sometimes present, sometimes absent, and

A perceptual state with internal characteristic *S* is a presentation of property *F* if and only if its function is to *detect* the presence of a thing with property *F*.

This schema is open to an objection that we must now attempt to overcome. The objection is based on the view that epistemic and practical assessment feed into and modify perception as much as perception feeds into assessment. Nelson Goodman puts the point with characteristic panache:

The Kantian dictum echoes here: the innocent eye is blind and the virgin mind empty. Moreover, what has been received [i.e., by perception] and what has been done to it [by reason] cannot be distinguished within the final product [belief]. Content cannot be extracted by peeling off layers of comment.¹⁹

If Goodman is right, it is not possible to speak of perception as detecting anything. We have beliefs about what is the case and no other epistemically significant states that constitute, as it were, pre-beliefs. So to speak of the *content* of what perception contributes to belief and action is an impermissible abstraction.

It may very well be right to say, as Goodman does, that perceptual content cannot be got at by peeling off layers of judgmental content. But I think that it is possible, nevertheless, to get at what perception contributes to judgment, simply because it is possible to create situations in which the output of perception gets isolated by the fact that judgment is not allowed sufficient information with which to comment on it. Techniques for achieving this result experimentally are in evidence in the pioneering work by Bela Julesz and Land.²⁰

Julesz was dealing with depth perception by means of measurements of binocular disparity, i.e., the displacement of an image of something on one retina with respect to the other. The big problem that has to be solved before this sort of depth perception can fully be elucidated is this: How do we manage to identify two images as of the same object—how do we *match* images in order to measure their binocular disparity? Some psychologists of perception have tried to solve this problem at a very high level; they proposed that the world is first segregated into big external *objects* (such as geometrical figures on a ground or even objects such as pieces of furniture or human beings) by means of some innate (so says the Gestalt theory) or learned ("New Look" theories) process of identification, and that the images of these objects are then matched. This is the type of proposal which leads to Goodman's sort of skepticism about the separability of perception and other mental processes, for it makes

¹⁹ *Languages of Art*, 2nd ed. (Indianapolis, IN: Hackett, 1975), p. 8.

²⁰ Julesz, "Binocular Depth Perception of Computer Generated Patterns," *Bell Systems Technical Journal*, xxxix (1960): 1125-1162; and Land and McCann, "Lightness and Retinex Theory," *op. cit.*

judgment influence perception just as much as perception influences judgment. If this sort of theory is right, it would indeed be wrong to identify the output of perception as the input to judgment.

Julesz disproved the contention that the measurement of binocular disparity is a process that operates at a "high level" by recreating depth perception on "random-dot stereograms." These stereograms consist of two images: one is a computer-generated random pattern of dots, and the other is a repetition of the same with a square region shifted over by a few minutes of arc. When these images are presented, one to one eye and the other to the other, it appears that the square region is closer to the observer than the rest. This shows that the binocular disparity of the square region was detectable. But, because the pattern consists only of randomly placed dots, the measurement could not have needed prior segregation of objects within the image.

Land's experiments concerned the perception of the brightness of an object. (When the brightness perception from all color channels is combined, this sort of perception amounts to color perception.) We have seen that the perceived brightness of an object exhibits a kind of constancy not displayed by the light it reflects under varied circumstances—for example, a uniformly colored wall will look uniformly colored, even when it is half in shadow. Now, this is a natural place for a high-level theory. How are we to explain the fact that different light information culminates in the same appearance? Why not say that our belief that the wall is uniformly colored informs our perception of it as uniformly colored? Land disproved this supposition by showing that humans are able to detect color constancy even in random patchworks of color—"Mondrians" as he called them, because of their resemblance to the abstract works of that artist. The Mondrians are presented to observers who have never seen them before; and these observers are able to discriminate different colors that are so illuminated as to send the same signal to the eye, and identify patches of the same color even when they are so illuminated as to send different signals to the eye. The constancy phenomena are thus displayed in the complete absence of any beliefs about the actual brightnesses present. Moreover, the Mondrians lack objects with which colors or brightnesses are inductively associated—apples, leaves, etc. So another source of interference by beliefs is eliminated.

I conclude from these results that it is at least sometimes possible to isolate what is detected by perception independently of judgment, decision, and the like. Whether or not it is true that perception can be influenced by judgment, this suggests that what perception detects exists independently of judgment. We shall now pursue schema

P, the thesis that a perceptual state presents that which it is its function to detect.²¹

VIII. REALISM AND THE FUNCTIONAL SCHEMA

It is a consequence of schema P (the *functional schema*, as I shall call it) that the properties presented by perception exist independently of perception. Understood in the light of the functional schema, Land's theory is a hypothesis about the functional organization of color vision. The function of color vision is to detect surface color. The (supposed) activity of the visual system in detecting brightness "edges" in the image incident upon the retina and in processing this information in the way described earlier is subordinate to the function of detecting color. The fact that mediates this subordination is, simply, the correlation between brightness edges and color discontinuities. It was, in other words, the historical operation of the correlation, together with the advantage of possessing a color detector, which was responsible for the evolutionary persistence of this edge processing.

The claim that content can be attributed to perceptual states as a result of such hypotheses about functional organization has two immediate consequences. First, it provides a justification for expressing the content of color vision in distal terms. Suppose that somebody were to say that Land's theory shows not that we see color, but that what we have so far mistakenly supposed to be color vision is, in fact, a system for detecting different sorts of brightness gradients in the (proximal) image incident upon the retina. This hypothesis about content is perfectly compatible with everything that Land says about the mechanism of color vision as it now exists. Only when we take the evolutionary development of color vision into account do the two ways of describing Land's theory turn out not to be factually equivalent. The functional schema gives crucial importance to the (supposed) fact that the ability to distinguish different sorts of brightness gradient in an image does not by itself confer an evolutionary advantage upon us. Rather, this discriminatory ability persists only because, as it happens, it happens to correlate closely with the ability to discriminate surface color.

A second important consequence of functionally identifying the

²¹ In *The Modularity of Mind* (Cambridge, MA: MIT Press, 1983), Fodor argues that in fact judgment has a very limited influence on perception: there is no feedback to perception except when calls on memory are built into the computational routines that constitute perception. Similarly, Marr allows memory to influence perception only after what he calls the "2 ½ D sketch" has been formed. Both these conditions constitute very significant restrictions on the unlimited feedback that Goodman and others of his ilk posit. But I have tried to make my argument independent of these restrictions.

output of the system we are discussing as information about *color* is that it makes the system subject to what I earlier entitled normal misperception. As we saw in section III (and note 12), it is possible to create a surface on which color varies spatially in the manner normally characteristic of illumination gradients. Such a surface would, when uniformly illuminated, be perceived as uniformly colored. The functional schema enables us to identify this as a misperception—which is surely what it is.

Together, distality and the possibility of misperception constitute realism in the theory of perception. Distality implies that perception purports to be of features of the world which are independent of perception itself. That, of course, does not imply realism, since it does not imply the actual existence of the independent world. But add the possibility of misperception, and we now have realism; for misperception would not be possible if the world of perception did not transcend perception itself.

In being thus committed to realism in my interpretation of the new theories of perception, I find myself at odds with a number of philosophers—Keith Campbell, C. W. K. Mundle, Clyde Hardin, and James McGilvray—who have discussed computational theories of perception in general, and color vision in particular. They have two lines of argument in support of an antirealist interpretation.

The first line of argument, which is emphasized by Hardin and Campbell,²² assumes that a color, blue for example, will be defined as the real property that, in normal circumstances, produces a particular sort of quale in us. But, they say, there is no such property.

Campbell and Hardin give two sorts of examples in support of this view. The first kind is irrelevant. Hardin says: "Sapphire is blue because of the transference of electrons from iron to titanium, but lapis lazuli is blue because of the vibrational energy characteristics of conjugated bonds" (496; cf. Campbell, 137). So what? It does not follow that they have no relevant real property in common—for example, reflectance.

The second sort of example is more troubling, at first sight. Perceived color is not all due to reflectance. Some surfaces may be blue because of reflectance, but the sky is blue because of dispersion. There is no real property common to such diverse causes of color perception. I shall make two comments about this. In order to make the point that Hardin and Campbell want to make, one must show that misperception is not involved. Why should we not say that the

²² Campbell, "Colours," in Robert Brown and C. D. Rollins, eds., *Contemporary Philosophy in Australia* (New York: Humanities Press, 1969), pp. 133–157; and Hardin, "Are 'Scientific' Objects Coloured?," *Mind*, xciii, 372 (October 1984): 491–500.

sky is *misperceived* as blue? Maybe the evolutionary advantages of color vision pertain *only* to the discriminatory resources it affords with respect to surfaces. Perhaps perceiving the sky as blue is just an artifact of this mechanism and it neither increases nor decreases the fitness we get as a result. That would be the sort of case in which we should want to say that the sky is (normally) misperceived as blue. Thus, the fact that it is perceived as blue in normal circumstances does not imply that it *is* blue.

There is, however, a more general point in Campbell and Hardin's arguments which needs to be considered. We perceive not only reflecting surfaces as colored, but also translucent, iridescent, and luminous surfaces. The reflectances of these other surfaces is quite irrelevant to their perceived colors. Surely it makes no sense to say that these are all misperceived. But this still does not prove the conclusion. What it proves is that it is not *reflectance* that we perceived, but some extension of reflectance, some property that is coextensive with reflectance when restricted to reflecting surfaces. I shall not attempt to define this property here, but I believe that it would not be hard to define.

A *second line of argument* is emphasized by Mundle and McGilvray.²³ These philosophers argue that, according to Land, our awareness of what we call "color" is the product not only of the stimulations incident upon the retina, but of our own "mental" activity. Thus, they assimilate color vision to something like the Kantian model of transcendental idealism, arguing that, though color vision is initiated by a signal from without, the end product represents the external world in a way that owes more to the operation of color vision than to the nature of properties that have real existence.

The position taken by these philosophers is, I think, essentially the same as supposing that "color" vision is actually a presentation of certain kinds of brightness gradient. The claim is that there is no way to identify the output of color vision except as the input stimulation processed in a certain way. They do not question that this output should be *called* "color": their point is that, whatever it is called, there is no reason to identify that output as representing an external property. I claim, of course, that this is wrong: the functional schema—and the historical explanations it implicitly invokes—provides us with just this reason.

IX. INTENTIONAL EXPLANATION

A strength of the functional schema is that it accounts well for the peculiarities of intentional explanation—the explanation of the

²³ Mundle, *Perception: Facts and Theories* (New York: Oxford, 1971): chs. 9 and 11; and McGilvray, "To Color," *Synthese*, LIV, 1 (January 1983): 37–70.

causal interrelations of states by reference to the content assigned to them.

The first point to make here is that the functional schema enables us to distinguish between states that should be assigned content from states that should not. In this respect, the causal theory is not helpful. Recall that, according to a causal theory, the content of a perceptual state is what normally causes states of that material type. There is nothing in this theory itself which prevents us from saying that, for example, a column of smoke is a perception of wind speed, since it has a feature—the angle it makes with the horizon—that is normally caused by a particular wind speed.²⁴ Thus, a causal theory needs an adjunct that distinguishes content-bearing states from non-content-bearing states. This does not apply to the functional schemata, Q and P. These assign content only to states that have the function of registering the presence of some object or circumstance. Since the angles of smoke columns do not have functions, they are excluded from the purview of this schema.

The second point to make is that the functional schema helps us understand the *parallel relationships* employed in intentional explanation. In intentional explanation, the explanantia and explananda are states with content. The explanation presupposes that the explanantia states cause the explanandum state to come into being²⁵; the explanatory weight, however, lies not on this causal relationship, but on the logical relationship that exists between the intentional objects of the states. For example, suppose we seek to explain *X*'s belief that *p* in terms of his belief that *q* and his belief that if *q* then *p*. This explanation would be inadmissible if the explanantia beliefs had not caused the explanandum belief to come about. The force of the explanation, however, is not primarily due to this causal relationship, but is due to the logical relationship that holds between *q*, if *q* then *p*, and *p*. Once we are told the relationship of the *content* of *X*'s belief to the *content* of his other beliefs, we understand why he believes that *p*. It is this obliqueness of accounting for states by reference to content which characterizes intentional explanation.

It is puzzling that it should be legitimate, in explaining why *X* came to believe that *p*, to appeal to the logical relations that *p* bears to the objects of *X*'s other beliefs—on the face of it, this explanatory strategy involves a category mistake, a confusion between the propositions toward which beliefs are directed, and the beliefs themselves.

²⁴ The example is David Hills'; see "Mental Representations and Languages of Thought," in Ned Block, ed., *Readings in the Philosophy of Psychology*, vol. 2 (Cambridge, MA: Harvard, 1981), pp. 11–20.

²⁵ That intentional explanations require causation is, of course, a point rediscovered by Donald Davidson in "Actions, Reasons, and Causes," this JOURNAL, LX, 23 (Nov. 7, 1963): 685–700.

For, if we are interested in the transition from beliefs already held to new beliefs, we ought to attend to the causal powers of beliefs. Why should these causal powers be determined by abstract semantic relations in the realm of propositions? To answer this question is essentially to show why intentional explanation is legitimate. And, if a particular scheme of intentional ascription makes it easy to answer the question, that increases its plausibility as well as the plausibility of the explanations it supports.

It is easy to see that the functional schema helps explain why states should be connected in this way. If the function of a state is to detect the truth of p , and of another state to detect the truth of if p then q , then, if the detection of q is of any advantage at all, we should expect that a connection might well develop between these states and the state that detects q . By contrast, a causal theory is of no help here at all. That a state is normally caused by the truth of p , and another one by the truth of if p then q , is no reason at all to expect that the two together would cause the state that is normally caused by q . This is a consideration which applies not just to perceptual states but to all states with content—it suggests that they should all be defined functionally.

Evolutionary epistemology has sometimes been regarded as providing us with reasons for disavowing skepticism. It is easy to see now that it cannot do this. For evolution may provide us with senses that are quite reliable, but there is no reason to believe that it provides us with senses that guarantee the truth of any particular perception—that is the lesson of “normal misperception.” Nevertheless, evolutionary epistemology has an important role to play. We started this paper with a question from traditional epistemology—What constitutes the mode of presentation in perception? We found that computational theories of perception have something important to contribute to this question and that the theory of evolution is foundationally important to computational theories. To the extent that such theories tend to inform us about the structure of our perceptually derived concepts, then, evolution can after all do what Konrad Lorenz,²⁶ the founder of evolutionary epistemology, hoped. It can show how that which is a priori for the individual is a posteriori for the species. That is, a study of evolution can help us lay the foundations of an understanding of innate ways of representing the external world. That is no small contribution.

MOHAN MATTHEN

University of Alberta

²⁶ See “Kant’s Theory of the *A Priori* in the Light of Evolutionary Biology,” in Konrad Lorenz, R. I. Evans, trans. (New York: Harcourt, 1975), pp. 181–217.