

# Reply to Rescorla and Peacocke: Perceptual Content in Light of Perceptual Constancies and Biological Constraints

TYLER BURGE

*University of California, Los Angeles*

*Origins of Objectivity* makes, I think, four primary contributions. The first is an elaboration of perceptual anti-individualism—the view that the natures of perceptual states depend constitutively on relations, including non-representational causal relations, to an environment. The main innovation is to distinguish within non-representational causal relations an important subset of teleological relations—those (like eating) involved in fulfilling whole-individual functions. The causal relations, between individuals and their environment, that are involved in fulfilling these non-representational teleological relations help set a framework within which representational natures of perceptual states are determined. The teleology grounds a system of natural norms—standards for well-functioning. Like the non-representational functions, and the acts and processes that fulfill them, these norms are pre-representational. By reflecting on perceptual and higher-level capacities and *their* functions and norms in the context of pre-representational ones, we gain insight.

The second primary contribution is a criticism of a syndrome of views that dominated twentieth-century philosophy and that still grip some philosophers. These views hold that representation of the physical environment requires a propositional, or even linguistic, capacity to *represent* constitutive conditions for such representation. I charge all forms of this idea with hyper-intellectualization, with controverting common sense about animals and young children, and, most importantly, with contradicting scientific knowledge about perception. Perceptually attributing properties to environmental entities does not require any of the higher-level capacities invoked by those views.

The third main contribution—perhaps the most important—is an account of the lower border of representational mind. This is the border at which a distinctively psychological kind of representation begins. The relevant sort of representation sets veridicality conditions—conditions for accuracy or for truth—and is an aspect of the nature of a psychological state. I contrast this kind of representation with another phenomenon often also called ‘representation’—causally based, statistical correlation that has a function. I call this latter phenomenon ‘*information registration*’. Information registration occurs in plants and water pumps. It is not distinctively psychological. The most primitive, distinctively psychological type of representation is perception. The lower border of perception is the lower border of representational mind. I locate this border phylogenetically—in certain arthropods. Perception constitutes a primitive type of objectivity—accurate representation of a mind-independent, or constitutively non-perspectival, subject matter. I provide a constitutive account of perception. I take perception to be objective, sensory representation by an individual. I discuss each of these conditions. Regarding objectivity, I highlight perceptual constancies.

The fourth primary contribution is a discussion of scientific knowledge about some categories of perceptual representation—body representation, representation of aggregate ratios, tracking small numbers of particulars, spatial representation, and temporal representation. This discussion refines the distinction between representational states and information-registration states that also guide behavior—often in virtuosic, computationally complex ways. The discussion elaborates the general account of perception, and provides constitutive accounts of specific categories of perceptual states.

I am grateful to my commentators for their provocative comments and for engaging with the book.

I begin with Michael Rescorla’s discussion of representational contents in perceptual constancies. Three views figure in his discussion. The first is an ‘extreme Russellian view’ that dispenses with perceptual attributives and confines its account of perception to citing a relation between perceiver and physical attributes, and their instances. The second holds,

(II) If a perceptual constancy centers on an attribute, distinct exercises of the constancy, prompted by differing proximal inputs, commonly yield perceptual states with different representational contents.

The third holds,

(III) If a perceptual constancy centers on an attribute, distinct exercises of the constancy, prompted by differing proximal inputs, commonly yield perceptual states featuring distinct perceptual attributives that represent the same physical attribute.

Rescorla rejects the extreme Russellian view—as do I.<sup>1</sup> Rescorla seems to accept (II), calling it a weak (non-ambitious) claim. I accept it. Rescorla doubts (III). In *Origins* I argue for it.

Consider an exercise of shape constancy. One has different perceptions as of a specific rectangular shape, deriving from perceiving it at different orientations. (III) holds that, commonly, there are different perceptual attributives for the rectangular shape.

Rescorla mentions a ‘minimalist’ reading of (III). On this reading, the different perceptual attributives for shape are formally complex—like rectangular at tilt T and rectangular at tilt T<sub>1</sub>. Further, there is only one “non-complex”, or basic, attributive that is proprietary to the property, (specific) rectangularity.

Rescorla discusses this minimalist line only briefly. He makes two criticisms. One is that it must defend the view that each percept ‘employs the postulated complex attributive’, rather than attributing rectangularity and tilt ‘separately’. The other is that the minimalist line is

available to an extreme “Russellian” who individuates perceptual contents entirely through their *representata*. Thus, the minimalist view does not secure a significant theoretical role for modes of presentation over and above *representata*.

I think that neither criticism carries weight. The very idea of ‘separate’ perceptual attribution here makes little sense. Normally, one cannot see a rectangular surface as rectangular without seeing the rectangularity as being at a specific tilt. In many cases, the tilt attributive is within the scope of the rectangularity attributive. Scope priority can derive from various sources—for example, being the focus of attention, or being the locus for organizing information for guiding action or for memory. Even apart from the scope point, normal perceptual representation of tilt, in mammalian visual systems in normal situations, intuitively and *inevitably* is connected to and qualifies perceptual representation of surface shape.

The criticism that the minimalist line does not secure a significant role for modes of presentation against the Russellian view is also mistaken. Maintaining that there is a type of representational content over and above *representata*—a two-tier account—is secured through reflection on the science of perception. Computations invoked in the science operate (partly) on perceptual representational contents, never on environmental *representata*.

Perceptual constancies illustrate the key distinction between a perceiver’s ways of representing and the *representata*. This distinction is the crucial difference between Frege and Russell. The distinction does *not* hinge on whether

---

<sup>1</sup> *Origins of Objectivity*, 119–121, 362, 385–396; ‘Disjunctivism and Perceptual Psychology’, *Philosophical Topics* 33 (2005), 1–78.

there are different basic representational contents that represent the same *representatum*. In fact, although I do not agree, Frege almost surely thought that, in a perfect language of thought, there are *no such cases*.

The minimalist line is one elaboration of (III). My defense of (III) in *Origins* is compatible with it. The book neither accepts nor rejects it. Nothing in Rescorla's discussion counts against this elaboration. Hence nothing counts against (III). I do reject the minimalist elaboration. Before explaining this position, I address other issues.

Rescorla questions how close perceptual constancies are to Fregean informative identities. (*Fregean informative identities* are true identity thoughts that specify the same referent in different ways in the two argument places of the identity attributive.) Rescorla doubts that exercises of perceptual constancies attribute the relation *identity*. I agree. I think that perceptual representational content contains no logical constants. My characterizations of perceptual constancy may have misled. I have written of a capacity to perceptually represent some attribute or particular 'as the same' under significant differences in proximal stimulation. I did not take this formulation to require attribution of the relation *identity*. A better formulation is 'a capacity to perceptually represent some given attribute, or particular, as that attribute, or particular, under significant differences in proximal stimulation'.

Rescorla discusses passages in which I say that perceptual constancies are analogs of Fregean informative identities. I believe that his discussions are in some ways misleading. I will not detail these ways. Some of the problem lies in my not having discussed the semantics of perceptual constancies in detail.

There are several differences—beyond lack of an identity attributive—between exercises of perceptual constancies and Fregean informative identities. The primary one is that perception is iconic and non-propositional, whereas identity thoughts are propositional. But there are further differences. Different exercises differ in different ways. Let us consider an exercise of color constancy—one of Rescorla's examples—that elicits relatively clear analogies to Fregean informative identities.

Suppose that I veridically perceive instance *i* of color-shade *c* as being *c*, where the color-shade is illuminated by white light; and I continue—330 milliseconds later—to veridically perceive *i* as *c*, where *c* remains illuminated by white light. The types of perceptions formed at the two times are, let us suppose, the same. Compare that mundane situation with this exercise of color constancy. I veridically perceive an instance *i* of color-shade *c* as being *c*, where *c* is illuminated by white light; and I continue—330 milliseconds later—to veridically perceive *i* as being *c*, where *c* is illuminated by blue light. In this second case, the types of perceptions formed are different from one another. The ways of perceiving *i*—the perceptual attributives applied—are different. They differ even though they indicate and attribute the same color-shade.

Exercises of both pairings are fallible. (Note that even most Fregean identities of the form  $a=a$  are fallible, because of the possibility of reference failure.) But the second pairing is more informative than the mundane pairing. It constitutes a perceptual achievement that differs in *perceptual* information value from the first. Analogies to Frege's distinction between  $a=a$  identity thoughts and more informative identity thoughts seem clear.

We cannot account for the achievement by saying, as Rescorla suggests, 'the perceptual accomplishment consists in veridically applying a single perceptual attributive...when prompted by diverse proximal inputs'. Appealing to different proximal inputs does not explain the difference in *perceptual* states (II) and *perceptual* information values. The achievement cannot be explained in terms of a bare difference in demonstrative-like applications of the same attributives. For it is simply not true that I perceived the color-shade, in all *perceptual* respects, in the same way. The ability-general aspects of paired perceptual states differ. The only possible ability-general perceptual difference in the state-types is a difference in perceptual attributives that bear on the ways the color-shade occurs, (III).

More general considerations support taking the primary difference between the two pairs to be a difference in perceptual attributives that characterize the constant attribute.

A first consideration concerns accuracy conditions. Perceptual representational contents are conditions for perceptual accuracy. Conditions under which the two state-instances in the first pair are accurate are type-identical.<sup>2</sup> Conditions under which members of the second pair are accurate differ. If the perceptual system produced the same perceptual state-type in response to the blue-lit color-shade as it did to the white-lit color-shade, it would be inaccurate, even assuming that it got the color-shade right.

Imagine three drawings. Each is to represent, as accurately as possible, what the artist sees in looking at the color shade. In the first, the color-shade occurs in white light; the drawing is fully accurate. In the second, the color-shade occurs in blue light; the drawing exactly matches the first drawing. In the third, the color-shade occurs in blue light; the drawing represents the color-shade in a way that accounts for its being blue-lit. Clearly, the second drawing is not fully accurate—as the third may be—even though it represents the color-shade as the color-shade that it is. Analogously, accurately perceiving the color-shade requires representing the color-shade in a way appropriate to its being lit in the way that it is in fact lit. Representational contents set veridicality conditions. So the representational contents of the second and third perceptual states differ, (II). The relevant differences are perceptual attributives that represent the color-shade in a way appropriate to its lighting conditions, (III).

---

<sup>2</sup> By 'type identical' I mean that the same arrangement of attribute-types is indicated in the same way by the two members of the first pair of perceptual state-instances.

A second general consideration that supports recognizing different perceptual attributives in exercises of perceptual constancies concerns ways in which different perceptual state-instances enter into law-like representational psychological patterns. Perceptual representational contents help individuate perceptual states. They help constitute basic perceptual representational kinds. The representational differences between instances in the more informative pair of perceptions surely figure in psychological explanation. Pairings of the first, mundane type might be more likely to be accurate than pairings of the second type. The first pairing might be more likely to trigger an action faster than the second.<sup>3</sup>

So the representational contents in the more informative pair of instances differ from one another. The differences do not reside in occurrence-based applications alone. They reside in different ability-general accuracy conditions and are embedded in different law-like representational psychological patterns (associated with different perceptual information values). Since there are representational differences, there are different perceptual states, with different representational contents, (II). The differences concern how the color-shade instance is perceived. The only relevant, ability-general differences in representational contents are differences in perceptual attributives, (III).

As Rescorla notes, on my view, the different attributives must indicate and attribute the same “constant” attribute—here color-shade. (That is obviously so, on the minimalist elaboration of (III).) Contrary to what he suggests, I take it to be ‘logically guaranteed’—more precisely, *entailed by the content and semantics of the attributives*—that the different attributives exercised in a constancy indicate the same “constant” attribute-type. They do so even in cases of error. Differences in perceptual information value derive from differences in attributives that attribute a “constant” attribute.<sup>4</sup>

I return to the minimalist elaboration of (III). That elaboration holds that in all exercises of perceptual constancies that center on attributes, different perceptual states and representational contents in the different stimulus conditions involve different *canonical* perceptual attributives for attributes *other* than the attribute on which the constancy centers. These canonical attributives qualify a *canonical* attributive for that constant attribute. In the more informative color-shade pairing, *one* canonical basic attributive that indicates the color-shade is paired successively with each of two canonical basic attributives that

---

<sup>3</sup> Much of Rescorla’s discussion focuses on remarks that I make about *post*-perceptual behavior or mental activity. I think that his discussion fails to convey the point of such remarks. The remarks were intended to elicit psychological patterns that evince different representational contents of *perceptual* states.

<sup>4</sup> Two types of error are common in exercises of perceptual constancies: (1) attribution of an attribute or attribute-complex that a perceived particular does not have; (2) failure to perceive a relevant particular. (There are also rarer failures of attribute indication.) (1) covers a perceptual state’s attributing the correct color-shade, but getting wrong how it is lit.

indicate different illuminations. Combining color-shade and blue-lighting-illumination attributives in attributing the color-shade yields a complex attributive that attributes the color-shade *in a different way* from the way in which the shade is attributed if the canonical color-shade attributive were combined with a canonical white-lighting-illumination attributive. On this minimalist line, in a given modality (say, vision), there is always *only one* “basic”, proprietary canonical perceptual attributive for an attribute. I doubt this line.

One class of doubts centers on the fact that some law-like psychological differences in representational states do not seem to depend on different complexes of attributives. I cite some cases, in a tentative spirit.

Foveal perception of a particular as having an attribute is different in law-like intra-psychological representational ways, and in ways that bear on likelihood of accuracy, from peripheral perception of the particular as having that attribute. This fact supports the idea that the representational contents differ by having different perceptual attributives. The differences do not depend on the complexity or non-complexity of the attributives. The point counts against believing that there is only one “basic” canonical attributive for an attribute. Some perceptual constancies depend on such perceptual differences. Consider location constancy with respect to an unmoved scene, with head-turnings or eye-shifts. Perceptual attributions of a given attribute can go in and out of foveal focus. Such attributional differences are embedded in different law-like patterns among representational states, and do not hinge *purely* on changes in representations of ancillary attributes.

Color constancy depends on sensing illumination differences. It is an empirically disputed issue whether different illuminations are *perceptually represented as such*. It has been argued that although registering illumination is a cue in color constancy, illumination is not perceived as such. Thus, illumination is sensorily registered, but commonly not perceived as illumination. If this view is correct, one cannot account for perceptual differences in attributing a given color-shade by invoking canonical perceptual attributives for different illuminations.<sup>5</sup>

All phenomenally conscious perceptual states are embedded in different intra-psychological, representational, law-like patterns—and probably different likelihoods of accuracy—from those that unconscious counterparts of those perceptual states are embedded in. A conscious perceptual attributive for any given attribute differs from any unconscious perceptual attributive for that attribute. The difference does not depend on the complexity or

---

<sup>5</sup> If this view is correct, accuracy conditions might be set in terms of a range of attributives all of which indicate and attribute the color-shade. Each attributive in the range would be accurately applied only if the color-shade is accurately attributed *and* the color-shade is illuminated in the way appropriate to that attributive. Such attributives would never attribute illumination *per se* to an instance of illumination. I shall discuss the semantics of such attributives elsewhere.

non-complexity of the attributives. Whether or not this point concerns perceptual constancies, it counts against the canonical attributives idea.

These points invite exploration. But the basic reason for rejecting the minimalist line is that it ignores the iconic nature of perceptual representation. Visual perception occurs in an ego-centrally anchored, spatial coordinate system. So does tactile perception and spatial hearing. Given certain conditions on the use of such a system, it is inevitable that there are multiple attributives for any given attribute or attribute-complex that is attributed to a perceived instance of that attribute or attribute-complex. There are three conditions.

- (1) The ego-centric coordinate system is fixed to the perceiver—for example, to the cyclopeian eye.<sup>6</sup>
- (2) The relevant spatial attributes are represented by means of the coordinate system.

Thus, an attribute is not represented through an attributive—such as the linguistic predicate ‘edge’—that perhaps places the indicated attribute in the coordinate system, but without using the system in representing the attribute. To meet condition (2), an edge of a cube must be specified not merely as an edge, but through specifications that plot the edge in the coordinate system—give its length, shape, and orientation, using spatial specifications within the coordinate system.

- (3) There can be change in the spatial relation between the coordinate system and represented attribute-instances—either because the observer’s angle of view changes or because the environmental entities change.

All systems of spatial perception meet these three conditions. Given that they are met, then for any given spatial attribute or attribute-complex for a purportedly perceived instance, the perceptual system has multiple perceptual attributives that can represent it.<sup>7</sup>

---

<sup>6</sup> There may be more than one visual coordinate system for a given perceiver—for example, in mammals, one for the cyclopeian eye and one for each eye. I assume here a single system.

<sup>7</sup> There is plausibly a level of perceptual attribution that is canonical. For example, it follows from the contents of all the different perceptual attributives for an edge of a given length—that each represents the same length. This equivalence gives rise to a derivative attributive for the length. Such derivative attributives are plausibly canonical, and are—in the case of spatial attributives—plausibly amodal. They mark more generic psychological states. The amodal cases may ground coordination between different modalities of sense perception. But the perspectively different ways of representing a given length in different applications of the ego-centrally anchored coordinate system, within a sense modality, are psychologically and representationally basic. The generalization in the text applies to basic, perceptual, space-based attributives—non-derivative attributives for edges, 2- and 3-dimensional shapes, lengths, sizes, textural units, bodies, and so on.



Let me illustrate this point concretely. Visual perceptual representation has a structure relevantly like that of pictorial representation—supplemented in various ways. Take a visual representation of an environmental edge—the border of a surface. Imagine that the representation has the form of a drawing of the edge. The drawing is a line. Such a drawing would represent not only the edge, but also its shape, length, and orientation—all as such. As intimated earlier regarding “non-separate” representation, the minimalist elaboration of (III) is right in holding that, at least in a large class of cases, perceptual attribution of any one attribute is inevitably qualified by attribution of others. Just as one cannot draw a line without drawing its length, shape, and orientation, one cannot visually represent an environmental edge as such without representing its length, shape, and orientation, as such.

Suppose that one accurately visually perceives, as such, a straight edge of a surface oriented up-down in the frontal plane. Imagine a realist drawing of the edge. It is straight, up-down, and a certain length. Now consider change in the edge’s orientation. Suppose that one accurately perceives, in a perceptual constancy, the same edge, as a straight edge of the given length, tilted back at a 30-degree angle. Suppose that, to cue the difference in tilt, one’s perceptual system utilizes differences in the geometry of the retinal image projected from the textural elements on the surface that the edge borders (355–359). One perceives the textural elements as having the spatial arrangement and sizes that one perceived them as having when they were not tilted.

Imagine drawing the tilted edge realistically. The drawing of the edge and its length is different from the drawing of the same edge and length in the frontal plane. The line is shorter. The line representing both edge and length is a different way of representing the same attributes.<sup>8</sup> Similarly, the drawings of the constant attributes of the textural elements differ from the drawings that represent the same attributes when their surface is not tilted. The drawings of textural elements at the surface’s top, titled backwards, are smaller and closer together.

The ways in which the edge and its length, and the spatial attributes of the textural elements, are represented differ from the ways in which each of those attributes is represented in the first drawing. The different lines represent the edge and its constant length in different ways. The drawings represent the textural elements as having given (constant) sizes, shapes, and spatial inter-relations; they do so in different ways.

The drawings meet the three conditions. They are within an ego-centrally anchored coordinate system fixed to the observer. They use that

---

<sup>8</sup> Although in the present case, the line-representations of straightness are straight in both orientations, the line-representations of straightness might also differ—for example, if, in one case, the edge were occluded.

system in representing the edge, its shape, and its length, and in representing the textural elements. And the spatial relation between coordinate system and environment changes.

Perceptual representations, including perceptual attributives for these attributes, are relevantly similar to drawings. The indicated and attributed attributes that remain constant are each represented by modes of representation—attributives—that differ in the different stimulus conditions. The idea of a single canonical representation for each of the constant attributes is mistaken.

In this case, as in many others—though, as noted regarding color constancy, perhaps not *all* others—the perceptual constancy involves perceptual representation (as such) of a “non-constant” attribute in different stimulus conditions. *However, in this case, as in many others, the attribute that differs is cued, as differing, by the sensory differences that yield differences in the attributives that represent the \*constant\* attribute or attributes. The attributives that are most proprietary to the constant attributes represent them in different ways in the different stimulus conditions.* Here the difference in tilt is cued by the retinal differences that yield different ways in which the constantly sized and arranged surface-textural elements are represented.

The different attributives that represent the constant attribute or attributes are ability-general. They are embedded in different psychological law-like patterns—different likelihoods for accuracy, different relations with other representational states. These differences signal differences in attributive representational content in perceiving constant attributes. The most proprietary, most basic, perceptual attributives for each of the constant attributes—edge, edge-length, textural-element sizes and shapes, and spatial relations among textural elements—differ in the different stimulus conditions. The point made for representation of environmental edges applies equally to surfaces, three-dimensional shapes, bodies, and so on.

Perceptual psychology assumes an iconic format for perceptual representation. Phenomenology may point some toward this conclusion. The basic reason for it is, however, theoretical. The format of visual representation takes on some of the geometry of the mapping of light on the retina. Representation via ego-centrally anchored spatial coordinate systems owes much to the spatial layout of light registration by retinal receptors. Some of this layout of pre-perceptual registration of retinal information is preserved and co-opted by the perception-formation process. Of course, both the pre-perceptual sensory layout and perceptual content with 3-dimensional significance derive from mapping the projected spatial structure of the environment. The sensory layout is given 3-dimensional spatial content in the format of an ego-centrally anchored spatial coordinate system, at stages of computation beyond retinal registration. This format informs nearly all perceptual representation. Given that relations between a perceiver and the spatial environment change, this

format *necessitates* systematic, fine-grained differences in the ways in which any given spatial attribute of purportedly perceived instances can be perceptually represented. The differences occur in maximally non-complex, maximally proprietary perceptual attributives.<sup>9</sup>

The minimalist appeal to one canonical specification of each attribute, and one canonical specification of each attribute-complex, misses the egocentrically anchored, co-ordinate-system-informed nature of perceptual representation.<sup>10</sup> Some systems of language and thought utilize canonical representation. Canonical representation even has a place in perception—in certain abstract aspects of perceptual representation (see note 7). It has no place in basic perceptual attribution.

The iconic nature of perceptual content renders comparisons to non-perceptual beliefs irrelevant. In supporting his apparent belief in canonical perceptual attributives, Rescorla notes that there are many routes for forming the belief there is mercury in the beaker. Despite major variations in stimulus conditions, we can use a single concept mercury.<sup>11</sup> The example shows how different such beliefs are from perceptions. A fundamental difference between perception and non-perceptual thought is that the latter is largely *non-iconic*. It abstracts from concrete details that are constitutive to perceptual representation.

One can think about, or *conceive*, an edge as such without conceiving its length and shape. One cannot *perceive* an edge as such without perceiving it as having specific determinations of those attributes. These connections among attributive-types derive from the iconic, perspectival form of perceptual representation. That form grounds the fine-grained character of perceptual attributives.

Chris Peacocke and I agree on methodology and on many features of perception. I join him here in focusing on disagreement.

Most of his comment concerns an opposition between what he calls ‘the Biological-Constitutive view’ and his ‘Action-Answerability view’. The discussion rests on a very serious misreading of my position on relations between biology and perceptual psychology.

---

<sup>9</sup> The very notion of non-complexity needs qualification in reflecting on perceptual representation. I will develop these matters elsewhere.

<sup>10</sup> Analogous points apply for perceptual attribution of color, lightness, and motion. As applied to color and lightness, the points are not grounded in the ways coordinate systems work, but in the ways representation of these properties depend iconically on light intensity.

<sup>11</sup> I have at least two non-complex concepts that apply exactly to mercury—expressed respectively by ‘mercury’ and ‘quicksilver’. But the range of basic attributives that apply to a given attribute is vastly more limited in thought than in perception, unless thought incorporates perceptual content.

Peacocke accurately quotes passages in which I maintain that there is a constitutive relation between biological function and perceptual representational content. Biologically basic actions and other whole-animal responses to biological needs figure constitutively in background conditions for perception: perceptual states depend for their representational content partly on the role that perception plays in causing realizations of individual biological functions (292, 370–371). I qualify these claims by allowing perception by artifacts that have analogs of biological functions, perhaps artificial life (321). What is central is that perceptual content is constitutively constrained in a framework that takes account of whole-individual functions (369–376). The key idea is not organic life, but life-related teleology. Still, organic life underlies all known autonomous perceptual systems.

The life-related teleology that grounds a framework within which perceptual content is determined provides what I call a ‘first level of triangulation’ (271–273, 319, 324, 370–3, 413; cf. 70n). I wrote, ‘This first-level of triangulation does not suffice to ground objective representation. It is, however, a necessary framework for such grounding. It privileges environmental macro-entities as candidates for being objects of representation, *if* explanation in terms of representation is justified’ (273). This level determines no specific content.

A second level of triangulation determines specific content. The second level consists in capacities and relations that yield perceptual constancies. The specific content of perceptual states is determined by what among ordinary distal environmental causes the individual can sensorily discriminate, and by the modes of such discrimination, within perceptual constancies (272–275, 324–326, 342–347, 377–378, 407–408, 413). The first level of triangulation helps motivate standard construals of perceptual constancies, and protects them against wholesale, non-standard alternatives (sense-data, aggregates of undetached body parts, proximal stimuli). Such alternatives do not figure appropriately in zoological and ethological accounts of pre-representational fulfillments of basic animal functions.

Peacocke seriously over-interprets my postulation of a constitutive relation between biological function and perceptual content. He interprets the relation as “*generating*” specific representational contents, if a content can be associated with some biological function. Not one of the passages Peacocke quotes supports this interpretation. One will search the book in vain for such support—much less for claims that entail that contents like containing protein are generated. The second level of triangulation excludes this particular content on standard grounds: perceptual systems cannot discriminate protein, and relevant containment relations, within perceptual constancies. As noted, the constitutive biological-function/perceptual-content relations that I postulate generate *no* specific representational contents (273).

Similarly, nothing in the book *rules out* a specific perceptual content for an attribute *within* broad ranges of biologically useful attributes (such as ruling out an attributive for symmetry within the range of shapes), on the ground that the content never figured in an individual's carrying out a biological function. In fact, I specifically denied such a principle. I wrote, 'I do *not* hold that for each perceptual state there is some specific need or activity that is distinctive to its representational content' (320).

The constitutive relations between biology and perceptual representational content impose only a *generic* constraint on perceptual content. 'Realizations of animal functions constrain perceptual content' is no more a universal quantification than 'doctors heal' is. It is a statement of a deep generic connection that need not hold in every case. The constitutive relations entail that some perceptual representational contents figure in psychological explanations of life-related actions or reactions—not that every perceptual content does so.<sup>12</sup>

Connection to a biological function is neither sufficient nor necessary for a perceptual state to have a specific representational content. Peacocke's criticisms in his sections on over- and under-generation are simply irrelevant to my position.

There *is* the difference that Peacocke cites in our answers to Quine's inscrutability and indeterminacy theses. I think Quine's position mistaken because it postulates, as equally best with the standard ones, whole ranges of *representata* that do not figure centrally in biological-ecological explanations of whole-animal behavior—whereas the standard ranges do figure centrally in such explanations.

Peacocke follows Evans in holding that Quine's non-standard referents are unacceptable because they do not account for counterfactual sensitivity to boundaries. I think that the Evans-Peacocke claim underestimates Quine's position. That position allows for, indeed requires, behavioral sensitivity to boundaries of aggregates of undetached parts, as well as to boundaries of particular undetached parts of those aggregates. The boundaries of relevant

---

<sup>12</sup> There are many indications in *Origins* that the constitutive constraint is generic, not a universal generalization. There is the denial, just quoted (320), that every perceptual content must be associated with fulfilling some biological function. I write of constitutive relations between biology and perception as setting a 'framework' (271, 320–321, 407) or 'context' (321, 407) within which discriminative abilities and perceptual constancies are determined. I write that perceptual psychology is 'broadly' (320) integrated with zoology and ethology—to block the idea of a lock-step connection. I use the constraint *only* to explain generic points about perceptual content—such as that it concerns distal environmental causes (371, 407) and that it is incompatible with systems of non-standard environmental entities that are never mentioned in biological explanations (213–215, 321–322). I call the constraint an 'empirically grounded default presumption' (272–273). I never formulate it as the universal generalization that Peacocke's criticisms target.

aggregates coincide with the boundaries of bodies—say, bodies with rabbit-like shapes.

Peacocke's reply is that this defense overlooks an asymmetry. He claims that the non-standard view must refer to the rabbit-like shape of the (aggregate of) undetached parts, whereas the standard view need not refer to undetached body parts or their shapes. He writes, 'Any mention of the undetached rabbit parts can be eliminated...'. He concludes that there is an excess of interpretation—a postulation of 'explanatorily irrelevant distinctions', on Quine's non-standard alternative.

I think this asymmetry claim clearly mistaken. Any account of perception of bodies as such must refer to undetached body parts—the textures and surfaces of bodies and the parts of bodies perceived as such under occlusions. One perceives whole bodies *only* by perceiving their undetached parts. Standard and non-standard views must refer both to entities with three-dimensional shapes that, say, rabbits in fact have and to those entities' undetached parts. The standard view appeals to bodies and body parts. The non-standard view appeals rather to aggregates of undetached aggregate parts, and aggregate parts.

I think that explanation of whole-individual teleological responses to the environment grounds answering Quine. The answer holds-fixed explanations in the biological sciences. Quine postulated a fundamental divide between natural and putative representational sciences. He believed that an indeterminacy infects psychological sciences that is “over and above” the “indeterminacy” that resides in natural sciences—what he counted normal inductive indeterminacy. Neither Quine nor I need deal with any problem about induction in the natural sciences.<sup>13</sup> My point was that if, with Quine, one holds biology fixed, the minimal ways in which perceptual explanations must “mesh” (319) with biological explanations of animals' non-representational responses to the environment block the wholesale mismatches that Quine's non-standard accounts of language and perception invoke.

Bodies, more systematically than aggregates of undetached parts, are referenced in standard biological accounts of individuals' non-representational responses to the environment—eating, mating, navigating, and so on. Quine postulated an indeterminacy in representation that occurs even as standard biological accounts are held fixed. Dropping bodies from an account of perceptual representation by individuals that sensorily discriminate and track bodies would unhitch perceptual psychology from biological accounts of animal function. Perceptual psychology is not thus unhitched. So the non-

---

<sup>13</sup> Peacocke suggests, note 1, that a generalization of his view might solve the induction problem for biology. The suggestion is not specified sufficiently to be evaluated. Solving the induction problem for biology would, in any case, not be directly relevant to Quine's position.

standard accounts are mistaken. Perceptual psychology and biology are different sciences. They are, however, connected in fundamental ways that Quine overlooked.

Peacocke's Action-Answerability view competes with the second level of triangulation, not, as he thinks, the first. According to the Action-Answerability view,

(A-A) The correct attribution of representational content to perceptual states is constitutively answerable to the range of properties of actual and counterfactually possible actions of their subject that those perceptual states are capable of explaining (in combination with other states).

I think (A-A) mistaken.

In the practice of perceptual psychology, specific perceptual content is set by discriminations in response to the environment, discriminations embedded in perceptual constancies. This method is refined by sources of fine-grain discussed earlier. The intra-psychological, law-like patterns that exhibit fine-grain may include, but do not depend on, patterns that incorporate actional psychological states. However, the *attributes* that perceptual attributives indicate are entirely fixed by discrimination embedded in perceptual constancies. The discrimination may include perceptual anticipation and perceptual memory. (As noted, the pre-representational teleological framework motivates standard construals of perceptual constancies and protects them against non-standard re-construals.) Action figures evidentially, but not constitutively, in specific perceptual content determination. Perception is *very* closely linked with action. But it is not constitutively impossible for a perceptual system to make more or fewer perceptual discriminations than can be reflected in action by a perceiver of the relevant type.

(A-A) is incompatible with scientific practice. Perceptual psychology does not hold its attributions of specific content answerable to accounts of action, or any other whole-individual uptakes from perception, perceptual anticipation, or perceptual memory. There is no armchair requirement, justifiable in opposition to the science, that every sensory discrimination that enters into perceptual content must be associated with a capacity for action—or even a capacity for *either* active *or* non-active whole-individual behavior—that uses that content. (A-A) borrows too much from behaviorism and is at odds with the science.

I turn to questions that Peacocke raises about perception of mechanical relations. He asserts that representational contents of some human perceptions have entailments about force—contents like leans on. I believe that this assertion is probably right. The issue is, however, debated in mainstream perceptual psychology, and not decisively settled. The issue cannot be settled from the armchair. The problem lies in determining whether the

most primitive representation of such relations is contributed by a perceptual system, or only, on each occasion of application, by a combination of the products of a perceptual system with those of long-term perceptual memory or those of higher cognitive systems. The methodological issues are complex. I will discuss them elsewhere.

One visually perceives mechanical/causal relations only by perceiving spatio-temporal configurations. If representations of mechanical relations are perceptual, as I expect some to be, they must mark different discriminatory abilities from those marked by perceptual representations of the kinematical (purely spatio-temporal) configurations on which the representations of mechanical relations are grounded. In *Origins*, I made a similar point in differentiating visual representations of *body* from visual representations of the generic shape that is the main recognitional basis for perceptually attributing *body* (458–461, 469).

Peacocke cites a discriminatory difference between a perceptual state that represents the spatio-temporal configurations that recognitionally signal *leaning* and a perceptual state that represents *leaning*. The latter, but not the former, is associated with a perceptual anticipation that would be unfulfilled, if a supporting body were removed and the other body did not fall. I agree that this discriminatory difference is necessary to distinguishing the mechanical from the kinematic perceptual representation.

Peacocke claims that even if such anticipation were present, *that* would not be enough for *leaning* to be indicated by perceptual content. He requires ‘further connection with contents concerning mass or force’. He does not justify this claim. I am not persuaded.<sup>14</sup> If a relation to *leaning* is the best explanation, among environmentally relevant alternatives (469), of the grouping of the relevant spatio-temporal patterns among bodies, and of the association of that grouping with the relevant anticipation, then the anticipation would suffice to differentiate perceptual representation of *leaning* from that of the spatio-temporal configuration. I doubt that, to perceptually represent *leaning*, an individual must be capable of richer representational connection to force.

I reject, for example, armchair requirements that an individual must have a primitive folk theory of force, or some non-conceptual analog. Peacocke and others have made such claims elsewhere. One must have capacities that differentiate *leaning* from environmentally relevant alternatives. One must track *leaning*, not merely associated spatio-temporal configurations. The

---

<sup>14</sup> I focus on force. Perceptual representation of *leaning* may require perceiving the participants as bodies. Visual body representation, however, is subject to quite spare conditions (454–470). It does not, I think, constitutively require representation of mass, weight, or force. Perceptual representation of various other types of causation does not require perception of causal participants as bodies.



tracking need not, however, involve ‘further connection’ to the difference, beyond the anticipational difference. Perceptual psychology’s practice in investigating primitive representation of causal relations, in Michotte’s and DeBiran’s early work and in current refinements of that work, imposes no such requirement. I think that no argument against the scientific practice will succeed.

I appreciate Peacocke’s suggestion that perceptual constancies might help ground understanding of higher-level representational capacities. I am more sceptical of safety accounts of knowledge than he appears to be. I think him right to suggest that philosophical accounts of empirical cognition and epistemology would be enriched by relying on an account of perception closer to scientific understanding than traditional accounts have been. My current work aims at understanding the upper border of perception, and relations between perception and those cognitive capacities, including perceptual belief, that immediately utilize it.

Both commentators share my sense of the importance of science in this area of philosophy. Only by doing philosophy of perception that makes serious use of science, and only by connecting normative considerations of epistemology and action theory with psychological states that individuals really have, will philosophy make progress in these areas. Philosophy of perception must become largely a branch of philosophy of science. Since it must, it will.