

# Origins of Perception\*

**Tyler Burge**

University of California, Los Angeles

In these lectures I reflect on powers of mind. Traditionally, salient generic powers have been taken to be four in number: sensory registration, perception, reason, and critical reason (or reflective deliberation). There is, of course, dispute in the tradition. Some held that differences among these powers are mere “matters of degree” – degree of clarity and distinctness, for example. Collapsing reason into critical reason was common.

In the last half-century, some held that perception is just sensory registration under complex information processing. Some are deflationary about mentalistic talk altogether.

I believe that there are relatively clear distinctions among these four powers. In the case of the higher powers, this belief is a matter of informed conjecture. In the case of the lower powers – sensory registration and perception – , there are decisive scientific grounds for acknowledging the distinction. These are natural psychological kinds, with relatively clear, empirically supported demarcations between them.

I also believe that there is substantial scientific ground for demarcating a range of post-perceptual types of mental states and events (beyond perceptual anticipation and perceptual memory) that do not count as reason, but that figure in cognition. In particular, there are intermodal capacities that take input from various sensory-perceptual modalities (vision, touch, hearing, proprioception) and represent at more abstract, intermodal levels. These capacities are not propositional, but they are not perceptual either.<sup>1</sup> So I believe that there are

\* This article is an expansion of the first of two Petrus Hispanus Lectures, given in Lisbon in 2009. A later version was given as the first of four Nicod Lectures in Paris in 2010. I also gave versions of the lecture at University College London, New York University, and University of Kentucky. I have benefitted from comments by

*Disputatio*, Vol. IV, No. 29, November 2010

Received on 12 November 2010

more distinctions among large kinds of mental competence than the four that I began with.

In these lectures I focus on the second and third members of this quartet – perception and reason. I will give glancing attention to the complex territory between these two psychological kinds. My primary interest lies in the constitutive conditions for these powers of mind. Investigating constitutive conditions is usually aided by reflecting on what is empirically known about the powers. I will reflect on some of what is known about animal and infant representational capacities.

I begin with constitutive conditions on perception.<sup>2</sup> Perception is, I think, the most primitive form of representational mind. It is the most basic capacity that exhibits *representation* in a non-deflated sense of the term ‘representation’. This point is important. I will dwell on it.

The term ‘representation’ is used in several ways. I want to distinguish my use. Explicating the term will serve clarity. But it will also indicate, I think, that many other ways of using the term obscure boundaries of what I take to be the most important psychological generic kind, certainly the most important generic kind along with consciousness. *Representational* psychological states are those that have veridicality conditions *as an aspect of their natures – as an aspect of the fundamental explanation-grounding kinds that they instantiate*.

Of course, nearly any state can be treated instrumentally as if it had veridicality conditions. The DNA of an organism can be treated as more or less accurately representing the organism’s phenotype. The orbital state of a planet can be treated as accurately or inaccurately representing the normal trajectory of the planet around the sun. I take it that in the psychologies of some individuals, having veridicality conditions is an aspect of the *nature* of the states – the kind of states they are – not simply a status accorded those states for the purposes of someone who describes the states. Being representational is, I

audience members on all of these occasions, and from correspondence with Georges Rey.

<sup>1</sup> See Susan Carey, *The Origins of Concepts* (Oxford: Oxford University Press, 2009).

<sup>2</sup> Some of what I say here condenses what I say in “Perceptual Objectivity”, *The Philosophical Review* 118 (2009), 285-324; and in *Origins of Objectivity* (Oxford: Clarendon Press, 2010).

think, a real feature of a state, not a feature of a stance toward the states.<sup>3</sup> The instrumentalist position just sketched has, I think, few adherents now. It is incompatible with the usual realist view of science. Science invokes representation as a kind embedded in law-like patterns. So there is empirical reason to take it as a real kind in the world.

Another, more common, way of deflating the notion of representation is to understand it in information-theoretic terms. On this usage, a state “represents” something if that type of state is statistically correlated with it to some relevant degree. This information-theoretic conception is often filled out with further conditions. The state type might be required to be causally dependent on types of things it represents. And the causal or statistical correlation might be required to have a biological function. I will say that any instance of this family of relations is a case of *information registration*. Information registration is a scientifically fruitful family of notions. Many of the sensitivities of plants and animals have been illuminatingly explained in terms of information registration. Many biologists, psychologists, and philosophers use the term ‘representation’ very broadly to include all information registration.

There is nothing in itself wrong with using the term ‘representation’ in this way. But such usage adds nothing to explanations in statistical, causal, or functional terms. And it obscures the distinctively psychological kind, *representation* taken in a non-deflationary way.

The tradition just discussed takes representation, construed as information registration, not to be a distinctively psychological kind. The notion applies just as literally to the sensitivity of plants to light, the sensitivity of paramecia to chemical compounds, and the sensitivity of water pumps to liquid levels, as it does to perceptual states. This extreme breadth of application indicates that the tradition tends to gloss over differences between information registration and the sort of representation that has traditionally been thought to be involved in perception, belief, language, and reasoning.

It is not enough to respond that all these phenomena are instances of representation – that is, information registration – in this broad sense. That may be true. But there is a scientific difference between

<sup>3</sup> For the instrumentalist view, see Daniel C. Dennett, ‘Intentional Systems’, *The Journal of Philosophy* 68 (1971), 87-106; reprinted in *The Intentional Stance* (Cambridge, Mass.: MIT Press, 1989).

information registration and representation properly so-called. In genuine representation, veridicality conditions play an apparently ineliminable role *in actual scientific explanation*. In explanations that appeal to information registration, talk of veridicality conditions is an after-thought. Reference to veridicality conditions is no more central to scientific explanations of the states of plants and paramecia than to scientific explanation of planetary orbits. The key feature of representation – that it can be accurate/inaccurate or true/false – plays no role in the so-called representation involved in information registration. Notions of statistical co-variation, causation, and biological function do not entail any genuine notions of accuracy or truth.<sup>4</sup>

Veridicality conditions figure centrally in explanations in some parts of psychological science. Specifically, explanation of perception – the most advanced part of psychology – centers on explaining how we perceive particulars and their attributes in the environment, and under what conditions we fall into illusions. This form of explanation has been present in psychology since Helmholtz’s work in the late 19<sup>th</sup> century. But it has become the center-piece of mathematically rigorous, systematic work that has shown all the signs of a mature science since the early 1970s. The science of *visual* perception has become a more impressive science than many parts of biology. Since science is the best basis for judging what sorts of things there are in the world, there is powerful reason to believe that representation, involving veridicality conditions as a key and apparently ineliminable feature, is a basic psychological kind.

I believe that it can be shown that this type of representation is not reducible to information registration.<sup>5</sup> I will not go into detail here. But the idea is that the notions of accuracy and error cannot be captured by any of the notions central to information registration. In particular, the notion of biological function cannot be the basis for explaining accuracy or error. Success in realizing a biological function is basically a practical matter – fitness for survival. But truth, accuracy, falsity, and inaccuracy are not practical matters. In principle,

<sup>4</sup> Veridicality is a genus. Accuracy and propositional truth are species. Perceptual states are accurate or inaccurate but not (as I use the terms) true or false. There is empirical reason to think that perceptual states do not have propositional structure, but are nevertheless accurate or inaccurate.

<sup>5</sup> See *Origins of Objectivity*, *op. cit.*, chapter 8.

biological success could be correlated with inaccuracy, all the way down; biological failure could be correlated with accuracy, all the way down. Information registration is neither sufficient nor necessary for representation, in the sense of ‘representation’ that involves veridicality conditions. Representation that sets veridicality conditions and information registration are different theoretical notions with different explanatory potentials.

So the first point that I make about perception is that it is a type of representational state, with non-deflationary veridicality conditions as part of its nature and with potential for grounding explanation. Perception is not to be assimilated to sensory information registration.

This distinction is not just a philosophical point. It shows up in how states of organisms are explained in perceptual psychology. Not all functioning sensory states are perceptual states, although all perceptual states are sensory states. What lies at the basis of this distinction in kinds? So far, I have indicated that the distinction is *signaled* by whether explanation of a state makes essential reference to the state’s having veridicality conditions – whether a state is the sort whose accuracy or inaccuracy figures in its grounding of scientific explanation.

Can more be said about what it is about perceptual states that makes explanations that invoke veridicality conditions fruitful, not just impositions or after-thoughts? The key to the distinction between perception and sensory information registration lies in a certain type of *objectification* involved in the formation of perceptual states.

This objectification is formation of a state with a representational content that is *as of* a subject matter beyond idiosyncratic, proximal, or subjective features of the individual. The subject matter here is the physical environment – which is mind-independent or at least, as in the case of color, constitutively non-perspectival. The relevant objectification involves a certain removal from the local or idiosyncratic. I will explain more concretely wherein perception involves this sort of objectification.

I assume that perceptual states are partly type-individuated in terms of their representational contents – contents that represent entities in the physical environment. I will assume that direct realist views – views that try to dispense with such representational contents – do not accord with the science. Direct realist views, in this sense, try to dispense with modes of presentation or ways of representing

that are distinct from the entities represented. They attempt to understand perception purely as a relation between a perceiver and the environmental objects of perception, without postulating any states that are ways the perceiver perceives such objects.

I believe that such views constitute some of the oldest errors in the philosophy of perception. In the first place, for any given attribute there are indefinitely many perceptual ways of indicating it; the science is primarily about these ways, not about the attributes that are indicated. In the second place, direct realist views do not account well for cases in which no entity is perceived. Some of them postulate properties that are not instantiated in the context, but that are ‘perceived’ or that ‘confront the perceiver’. But perceived entities cause perception; properties as types do not cause anything; and in the absence of an instantiation of a property in the context of the perceiver’s perception, no object of perception causes the perception. Moreover, some perception is *as of* never-instantiated physical properties. I think that never-instantiated physical properties do not exist. To postulate them as *objects* of perception in cases of referential illusion is to mistakenly project psychological modes of representation (that fail to represent anything) into a perception-independent physical ontology. The old error is conflating object of perception with mode of perceiving. In the third place, certain aspects of representational contents – such as defocus – have no counterpart in the represented world, and do not even *seem* to have such a counterpart. Defocus is not blur represented in the world; it is an inherent aspect of the representation – an inevitable deficiency in the mode of representation – that is a central topic of scientific study. Views that postulate only a perceiver and a relation to something perceived cannot accommodate defocus in any natural way. I think that there is no reasonable alternative to theorizing about the representational natures of perceptual states. Representational content, a representational way of representing entities that has structure and that sets conditions for veridicality, is – along with perceptual mode (vision, hearing, and so on) – the principal aspect of the representational natures of perceptual states.<sup>6</sup>

<sup>6</sup> I discuss these matters in more detail, with further objections to particular versions of direct realism in *Origins of Objectivity*, *op. cit.*, 362-364, 384-396 ; ‘Disjunctivism and Perceptual Psychology’, *Philosophical Topics* 33 (2005), 1-78. Examples of direct realist views are John Campbell, *Reference and Consciousness*

Let me say a little about the basic shape of perceptual psychology. The central problem of perceptual psychology, paradigmatically *visual* psychology, is the *underdetermination problem*. The problem can be characterized briefly as follows.

Although visual perception is of and *as of* particulars and attributes that occur in the wider macro-environment with which individuals interact in their basic biological pursuits, the initial states of the perceptual system are registrations of proximal stimulation. Such registrations of light arrays impacting the retina are not perceptions. But they are all that the visual system has to start with.<sup>7</sup> It is known that different environmental conditions can produce the same (or not-discriminated) registrations of proximal stimulations. So in this sense proximal stimulations do not determine their environmental causal antecedents – the entities that are perceptually represented. Correspondingly, the registrations of proximal stimulation underdetermine perceptual states that are accurate or inaccurate with respect to the environmental, causal antecedents.<sup>8</sup>

So there are two types of underdetermination. The science of perceptual psychology concentrates on the second kind of underdetermination. But the second kind underlies and helps ground understanding of the first. The *underdetermination problem* is the problem of answering the following question. How are perceptual states of and as of the environment formed, given that the proximal stimulations to which the system has immediate causal access do not determine either the environmental entities that the perceptual states represent as being there or the perceptual states that do the representing?

(Oxford: Oxford University Press, 2002); Michael Martin, 'The Transparency of Experience', *Mind and Language* 17 (2002), 378-425; and Mark Johnston, 'Better than Mere Knowledge? The Function of Sensory Awareness' in T.S. Gendler and J. Hawthorne eds. *Perceptual Experience* (Oxford: Clarendon Press, 2006).

<sup>7</sup> I oversimplify. The visual system utilizes a wider range of input, not just stimuli of the retina. The underdetermination problem applies to the wider range as well.

<sup>8</sup> Underdetermination is a mathematical matter: It is logically and mathematically possible for the environmental causes of the registration of proximal stimulation (the causes that are potential objects of perception) to vary while the registration of proximal stimulation remains fixed. It is logically and mathematically possible for perceptual states to vary while a given registration of proximal stimulation remains fixed. In fact, these sorts of underdetermination are always not only logically and mathematically possible. They are also psycho-physically possible.

Here is an illustration of the underdetermination problem. The input from the environment is represented by an idealized image that gives the radiance at each location in the plane of optical sensors for each wavelength of light. This input is registered in the visual system, inevitably with some loss of information. The registration of the impacts of light intensities on retinal receptors degrades the idealized image primarily through defocus.<sup>9</sup> By testing the sensitivity and sampling characteristics of the sensors, the science can determine how any given proximal input is degraded or defocused in the initial registration of the proximal stimulation by light. This initial registration – the beginning of the psychological process – is commonly understood, mathematically, as a two-dimensional array. What the science takes as an initial registration can be adjusted for context- and task-dependency. But the description that I have given applies for most explanations in vision science.

This initial registration, or encoding, of proximal stimulation is transformed through a series of events in the visual system into perceptual representations of, and as of, entities in three-dimensional space. There is a determinate optical and geometrical solution to the question of how a three dimensional array projects onto a two-dimensional coding of that array. But there is no determinate mathematical solution to how the two-dimensional coding is transformed into a representation of, and as of, a three-dimensional scene. The retinal encodings, together with all further input from proximal stimulation, underdetermine even the *physically possible* environmental causes. Perceptual states sometimes accurately specify environmental attributes and refer to environmental particulars that have such attributes. So perceptual psychology must discover laws that govern how registrations of proximal stimulations cause visual perceptions. These formation laws are distinctively psychological.

The formation laws, and law-like patterns of processing, *privilege* certain possible environmental causes over others. The effect of the privileging is that the undeterdetermining proximal stimulation triggers a perceptual state that represents exactly one of the many

<sup>9</sup> Defocus remains a property of the perceptual states that are eventually formed through processing in the perceptual system. Defocus is a property of the representational state that simply cannot be construed in terms of elements in the world that the perceptual states is perceptually related to: it is an aspect of the mode of presentation, the way those elements are presented.



possible environmental causes that are optically compatible with the given proximal stimulation. The underdetermination of environmental causes by proximal registrations renders the formation of perceptual states subject to error. Illusions occur when abnormal environmental causes produce the same types of proximal stimulations that are produced by normal distal causes. These conditions are a topic of the science.<sup>10</sup> The underdetermination problem and basic explanatory scheme of the modern science were stated by Helmholtz. The approach has become rigorous, mature science since the 1970s with the advent of computer simulations.<sup>11</sup>

Let us return to objectification. Objectification resides in the ways perceptual systems overcome proximal stimulation's underdetermination of environmental *representata*, and sensory registration's underdetermination of perceptual representation of those *representata*. The perceptual system distinguishes patterns in the sensory registration that are likely to be adventitious, or idiosyncratic to the subject, from patterns that tend to correlate with specific attributes in the environment. In effect, the formation laws and law-like patterns systematically distinguish the merely proximal from the probably environmental. Such patterns constitute the relevant *objectification*. Specification of physical entities is distilled out from the individual's sensory registration.

Objectification hinges on distinguishing what concerns the individual's receptors and what concerns a receptor-independent reality – and doing so in a way that indicates *specific* environmental particulars, properties, relations, and kinds.

There are many *non-perceptual* sensory systems that involve complex processing of information and impressive adaptation to environmental conditions. What makes a system *perceptual* is not always greater complexity or virtuosity in processing, or greater environmental adaptation. What makes a system perceptual is the particular type of processing involved. It is processing that contrasts registration

<sup>10</sup> Most points made here occur in any mainstream textbook in visual psychology. See Stephen E. Palmer, *Vision Science, op. cit.*, 9-11, 18-24, 55-59; and Vicki Bruce and Patrick Green, *Visual Perception: Physiology, Psychology, and Ecology* (Hillsdale, New Jersey; Lawrence Erlbaum, 1985, 2004 4<sup>th</sup> edn.). The real science resides in journal articles.

<sup>11</sup> David Marr, *Vision* (San Francisco, W.H. Freeman and Company, 1982) produced an early consolidation of the maturing science.

of proximal stimulation from states with veridicality conditions that specify elements of the environment beyond the sensory receptors. Empirical psychology explains perception in terms of capacities and processes that exhibit this distinction in the formation of representational perceptual states.

This contrastive process is the primitive objectification distinctive of perception. Objectification lies in marking off states that are *as of* specific system-independent elements in the environment from states idiosyncratic or local to the perceiver. The process is carried out in many different ways, *each specific to some environmental entity*.

Objectification in perception is implemented by *perceptual constancies*. I believe that perceptual psychology's use of genuinely representational notions – its postulation of states capable of accuracy and inaccuracy – is fundamentally motivated by the need to explain perceptual constancies. Perceptual constancies mark the distinction between perception and sensory information registration. They ground and motivate the scientific distinction between perceptual representational states – those sensory states whose natures involve potential for accuracy or inaccuracy – and sensory states that register information, but do not engage in genuine representation.

Perceptual constancies are capacities systematically to represent a particular or an attribute as the same under significant variations in registration of proximal stimulation. A corollary of this explication is that perceptual constancies are capacities to represent a particular or an attribute as the same from different perceptual perspectives, produced by different proximal stimulation.

Numerous constancies occur in the perceptual systems of animals. *Shape constancy* is a capacity to perceive a given shape under various stimulus and perspectival conditions. A square pattern can be seen as square whether viewed head on or at an angle. *Location constancy* is a capacity to represent a perceived entity as at a given distance and direction, under various types of stimulation deriving from various types of entities perceived. The sonar systems of bats, whales, and dolphins and the visual systems of a numerous animals can localize objects, even if the objects in a given location are changed so as to produce very different proximal stimulation. *Lightness constancy* is the capacity to represent a given lightness – on the scale from black to white – as the same under various stimulus conditions, including different illuminations.

Perceptual constancies are marks of objectification. I think that a

sensory system is perceptual if and only if the system includes perceptual constancies. The constancies exhibit the capacity of perceptual systems to contrast differences in proximal stimulation with environmental entities that are crucial to the perceiver's dealing with its world.

Not all selectivity with respect to proximal stimulation marks a perceptual constancy. All sensory systems neglect or filter noise. All such systems suppress some and respond to other information, according to the information's usefulness. Selective responsiveness to aspects of stimulation is a product of conditioning. All organisms adapt by habituation or conditioning. All conditioning constitutes accommodation to pressures of the environment. Some non-perceptual sensory capacities are explainable only by reference to complex quantitative processing. The difference between perceptual constancies and non-perceptual sensory selectivity is not a matter of selectivity or quantitative complexity. It is a matter of the nature of the processing.

Perceptual constancies exhibit structure and system in the processing that relate specifically to – and are explainable in terms of – *specific* environmental attributes. It is, I think, this specificity that invites explanation in terms of veridicality conditions – genuine perceptual, representational states. Perceptual structures differ from the generalized weighting and averaging techniques that mark sophisticated non-perceptual sensory systems. Perceptual learning also differs from the serial, piecemeal, averaging adaptation to proximal stimulation in non-perceptual sensory systems.

These points are not definitional. Perceptual constancies are distinctively *perceptual* capacities. There is no question of defining perception in terms of perceptual constancies. Such a definition would be circular. The sufficiency ('if') direction of my conjecture that a sensory system is perceptual if and only if the system includes perceptual constancies is trivial. Let me clarify the point of the necessity ('only if') direction. I see it primarily as an empirical conjecture about actual perceptual systems.

Some sciences do explain sensory states in terms of representational contents with veridicality conditions, and some do not. The idea of the conjecture is that the central aspect of perceptual systems that makes it necessary to explain states in them in terms of representational contents with veridicality conditions is the presence of perceptual capacities inherent in perceptual constancies. Perceptual

constancies are individuals' capacities to specify a single environmental element (attribute or particular) under many stimulus conditions that correspond to many different perspectives on the element. My conjecture is that, in the actual world, in the absence of the specific capacity that is a perceptual constancy, a system's ability to connect sensorily with environmental entities can be adequately explained in terms that do not invoke representational contents with veridicality conditions. As a matter of scientific fact, explanations of many sensory capacities do not need to – and do not – invoke representational contents that set accuracy conditions. Glosses on the explanations in terms of the accuracy or inaccuracy of the sensings are afterthoughts, not central components in the scientific theory. In perceptual psychology, explanation of the formation of accurate and inaccurate representational states is the fundamental aim of the theory.

I believe that, as a matter of fact, this difference in explanatory strategy in the sciences of sensory capacities corresponds to a difference in sensory capacities between those that involve perceptual constancies and those that do not. Representational accuracy and perceptual constancy are natural psychological kinds acknowledged in science.

I incline toward the more committal view that if a capacity could not represent any environmental element under a variety of stimulus conditions, and lacked any perceptual constancy, then as a constitutive matter it would be non-perceptual. This conjecture is not based on conceptual analysis. It is based on a judgment of conditions under which sciences are led to explain perception formation in terms of kinds of states that are accurate or inaccurate with respect to specific elements in the environment.

'Perceptual constancy' is a theoretical term in the sciences. It applies to a natural psychological kind. Differentiating perceptual capacities from selective, adaptive capacities that are non-perceptual is a matter of empirical theory. Many sensory capacities are simply not scientifically explained in terms of the formation of perceptual states with representational content. Science explains such non-perceptual sensory capacities in terms of (a) proximal stimulus conditions, (b) principles governing weightings or other mathematical operations on registrations from different sensors, (c) statistically significant correlations between the organisms' sensory states and elements in the environment, and (d) biologically significant functional relations between the organism's sensory states and entities in the environment

that they correlate with. Although these types of explanations are often called 'representational', they are not distinctively psychological. They make no essential use of representational contents with veridicality conditions. They apply to plants and amoebae just as much as to animals with perceptual systems.<sup>12</sup> Such explanations suffice to explain very many types of sensory states and capacities in living creatures. Such explanations contrast dramatically with explanations of perception in terms of states with conditions for accuracy regarding the physical environment.

As emphasized four and five paragraphs back, non-perceptual sensory capacities allow for variations in stimulus conditions under which an organism senses an environmental entity. Stimulus conditions are never exactly the same from one encounter to another. There are filtrations for noise and allowances for a range of variations of stimulus conditions. There are computational transformations and averaging of stimulus registrations. I take it, however, that it is a scientifically established fact that some of these "constancies" over stimulus variations are not perceptual constancies.

Providing an exact general account of what the difference is would be an interesting project. But I shall not try to explain the natural kind, *perceptual constancy*, in other terms. I take it that the types of variations and types of constancies differ in non-perceptual and perceptual cases. I take it that the allowances for variation in stimulus conditions differ in the two cases. In the perceptual cases, different variations are relevant to different specific environmental attributes that are perceived, whereas in non-perceptual capacities variations and their limits tend to apply equally across sensings of different environmental attributes. In the perceptual case, reference to states with accuracy conditions anchors explanation. In the non-perceptual cases, basic explanation invokes no such states.

*Perceptual constancy* is a well-established natural psychological kind. As a matter of empirical fact, many sensory capacities lack perceptual constancies. The differences between perceptual constancies and other, non-perceptual forms of sensory selectivity that I have summarized are real. They have grounded quite a lot of mature scientific theory.

<sup>12</sup> For detailed discussion, see *Origins of Objectivity*, *op. cit.*, 292-307, 403-406, 487-496, 529-531.

The constancies do not depend on knowledge or conceptual understanding. They occur in simple animals, such as bees and spiders, that lack capacities for propositional belief. These animals can represent an environmental feature in different ways, under very different stimulus conditions. Perceptually different representational contents – which mark different modes of presentation and which constitute different perspectives caused by different proximal stimulations – represent the same environmental attribute or particular, in different ways. So a basic feature of representation – the difference between the entity represented and the way of representing it, the representational content – is integral to the perceptual constancies.

As I have intimated, many species exhibit perceptual constancies in various types of sense-perceptual systems. Some arthropods, and most reptiles, amphibians, fish, birds, and mammals have visual perception. Most of the spatial constancies occur in these visual systems. Color constancy is scattered through the animal kingdom, apparently depending on how central color is to the life of the species. Birds and bees tend to have it. Many mammals appear to lack it. Object constancy has been demonstrated in many birds and mammals. Various aspects of touch, proprioception, and hearing are perceptual, again in a wide variety of animals.<sup>13</sup>

Perceptual constancies seem to play little or no role in most sensory systems for olfaction and taste. This is an empirical issue. If perceptual constancies turn up in these systems, the systems have perceptual aspects. Although it is plausible that these systems are largely non-perceptual, I think it more useful to think in terms of aspects of sensory modalities as perceptual or not, than to label a given sensory modality as perceptual or non-perceptual. For example, the sensing of saltiness in taste seems to involve no perceptual constancies. The salty registration is sufficiently directly correlated with saltiness in foods on the tongue that no perspectival competencies are needed (which is not to say that the registration cannot occur without the saltiness in the food). But there are constancies closely associated with taste. One can often locate a taste on the tongue. This capacity to localize tastes in space seems to derive from interaction between taste and proprioception, which does have perceptual constancies in locating bodily sensations with respect to a standing body

<sup>13</sup> See *ibid*, 419-421.

image. Whether one counts this a perceptual capacity “in” the system for taste or not does not seem to matter.

Similarly, localization of pain is commonly connected to such proprioceptive capacities for localization. Although the feeling of the painful quality of a sensation does not itself seem to involve perceptual constancies, or perception, the localization of the pain often does.

If other perceptual constancies turned up in the olfactory, or taste, or pain sensing systems, then although the systems might be correctly termed ‘non-perceptual’ (or ‘largely non-perceptual’) on the whole – that is, in most of their aspects – they would have perceptual aspects. Even dominantly perceptual systems like visual perception inevitably have many non-perceptual aspects. So it seems to me better, and certainly more precise, to think of specific capacities within, or specific aspects of, sensory systems as perceptual or non-perceptual than to be wedded to construing whole systems as perceptual or non-perceptual.

To sharpen the distinction between perception and sophisticated non-perceptual sensory capacities, I will discuss some examples of non-perceptual sensory capacities. One of the most common bases for navigation in the animal world is *homing* or *beaconing*. A beacon is a signal, such as a chemical, or beam of light, or recurrent sound, that emanates from a source. Beacons can be sensed and used to home in on their origin if the sensory capacity can determine the direction from which they come. A sensory capacity can determine the direction of a beacon’s source if it can differentiate among the intensities of relevant proximal stimulations in different parts of its body. For example, if the capacity can determine that the light intensity stimulating receptors on the individual’s right side is greater than the intensity stimulating receptors on the individual’s left side, then the individual can turn in the direction of the side that receives the more intense stimulation. By serially turning toward relevantly positive stimulation, or away from aversive stimulation, an organism can put itself in advantageous positions.

Variants on this method of navigation exist in a huge variety of organisms. Bacteria, paramecia, moths engage in homing with respect to light or dark. Ants use chemical emanations from other ants to trigger homing. Crickets find mates by homing in on mating sound signals. Snakes find prey partly by homing on a heat source. Salmon find their way back to their molting site, often covering thousands of

miles in the ocean, by following a series of olfactory chemical cues that occur in layers in the oceans. Some sea turtles appear to use wind-borne cues, probably either olfactory chemical cues or sounds, to find an island where they lay eggs.<sup>14</sup>

Some of these capacities involve retention of a series of sequenced stimulations.<sup>15</sup> Some involve complex averaging of the stimulation intensities on banks of receptors located in different parts of the organism's body. Such sensory capacities function to orient the animal in complex, adaptive ways in the environment. But these are not perceptual capacities. They operate entirely on the intensity and bodily location of registrations of proximal stimulation. In explaining homing capacities, there is no need to postulate a modeling of conditions in the environment beyond the sensory surfaces. There is no need to postulate a state with accuracy conditions. And science does not do so. No objectifying contrastive operation specific to elements in the environment occurs. No perceptual constancies occur in such sensory capacities.

A more complex computational mechanism that also yields virtuosic navigational feats is path integration. Path integration is a computational transformation that yields an informational vector, correlated with a distance and a direction, from information that correlates with speed, direction, and time. An animal can navigate away from its home by a convoluted route and return home by a nearly straight line. Path integration does not occur in the simplest organisms. But it does occur in arthropods – ants, bees, and bugs – , and in birds and

<sup>14</sup> Peter B. Johnsen, 'Chemosensory Orientation Mechanisms of Fish', in D. Duvall, D. Müller-Schwarze, and R. M. Silverstein eds, *Chemical Signals in Vertebrates* (New York: Plenum, 1992); M.J. Weissburg, 'Chemo- and Methanosensory Orientation by Crustaceans in Laminar and Turbulent Flows: From Odor trails to Vortex Streets', in *Orientation and Communication in Arthropods*, M. Lehrer ed. (Basel, Birkhauser Verlag, 1997); Paolo Luschi, Susanne Åkesson, Annette C. Broderick, Fiona Glen, Brendan J. Godley, Floriano Papi, Graeme C. Hays, 'Testing the Navigational Abilities of Ocean Migrants: Displacement Experiments on Green Sea Turtles (*Chelonia mydas*)', *Behavioral Ecology and Sociobiology* 50 (2001), 528-534.

<sup>15</sup> Psychologically relevant retention does not require genuine representation—representation that sets veridicality conditions. Psychology is free with its use of the term 'memory', applying it even to muscle memory. I use 'memory' to apply only to genuinely representational states. Nonrepresentational systems of sensory information registration can, however, rely on complex, long-term retention of sensory effects, as the salmon case dramatically illustrates.



mammals.<sup>16</sup>

For path integration to be possible, the individual's sensory capacity must contain an accumulator mechanism that is equivalent to integrating (a) the individual's motion without rotation (translational motion) and (b) the individual's rotation, as it makes its journey. The sensory bases for the accumulator mechanism vary with different animals. There are capacities, for example, that accumulate the energy units involved in taking a step and that measure the stress on joints produced by torque from twisted joints when turns or rotations occur.

Path integration takes many forms. All forms are more complex than I am indicating here.<sup>17</sup> All path integration capacities require retention of sensory input and input-processing governed by complex quantitative principles. But there is nothing in path integration *per se*, or in many path integration capacities in actual animals, that involves perception, or any other type of representation, in the sense of 'representation' that I sketched.

Let us take the point about perception first. Although there can be perceptual input into the accumulator mechanism, there often is not. In the proprioceptive case that I described, the sensitivity to energy involved in taking steps, and in recording stresses on joints under rotation, involves no perceptual constancies. The registration of the stimulation suffices for the input into the accumulator mechanism. No processing contrasts registration of proximal input from some further condition that is an objectified object of perception. Although the sensory capacity functions to enable the animal to adapt to the spatial environment, it does not yield perceptual states with accuracy conditions that represent spatial properties.

Path integration capacities are certainly "representational" in the broad deflated sense that counts any functioning statistical co-

<sup>16</sup> Ken Cheng, 'Arthropod Navigation: Ants, Bees, Crabs, Spiders Finding Their Way' in Edward A. Wasserman and Thomas R. Zentall eds., *Comparative Cognition* (Oxford, Oxford University Press, 2006); S. Zill and E.A. Seyfarth, 'Exoskeletal Sensors for Walking', *Scientific American* 275 (1996), 70-74; Ariane S. Etienne and Kathryn J. Jeffery, 'Path Integration in Mammals', *Hippocampus* 14 (2004), 180-192; Peter Berthold, 'Spatiotemporal Aspects of Avian Long-distance Migration' in S. Healy ed. *Spatial Representation in Animals* (Oxford, Oxford University Press, 1998).

<sup>17</sup> For more detailed discussion, see my *Origins of Objectivity*, *op. cit.*, chapters 8-10.

variation as “representational”. But in many cases, there is no representation in the specific, psychologically distinctive sense that requires that accuracy conditions figure in the nature of the sensory state, or in the grounding of scientific explanation. Explanations of the sensory capacities involved in much path integration do not need to appeal to states with veridicality conditions – any more than they need to appeal to objectification or perceptual constancies. So representation in a sense that involves veridicality conditions in a non-trivial way appears to be absent from the sensory capacities that yield input into many path integration systems.

The quantitative processing of information registered in the sensory capacities that feed the accumulator mechanism cannot add spatial significance that is not already present in the sensory information being processed. So the whole path integration system involves no spatial representation in a non-deflated sense.

Path integration systems evolved because they enabled animals to navigate space. But a functioning isomorphism between the path integration system and space does not suffice to make veridicality conditions an aspect of the psychological states themselves. So representation – in a non-deflated sense – is not involved in path integration systems *per se*.

The foregoing points are illustrative. Many path integration systems do have *perceptual* sensory input that has spatial representational content. In such cases, the whole system becomes an organ of spatial representation. My point is that a system’s functional relevance to space, and its capacity to deliver sensory states that make navigation of space possible, do not entail that the system involves spatial perception or any other type of spatial representation (understood in a non-deflated sense). What signals perception and genuine representation is successful scientific explanation’s appeal to perceptual objectification and perceptual constancies and, more broadly, states with veridicality conditions. In many cases, explanations of sensory capacities involved in path integration do not make reference to such states. Explanation in terms of sensory registration, isomorphisms, and biological functions suffice. There is no reference to space in the states of many animals that show a remarkable capacity to negotiate space.

There are many further cases in which animals’ sensory capacities enable them to cope with the physical environment in impressive ways, but in which no perceptual or other representational capacity is

involved. I hope that the foregoing illustrations suffice to provide some sense of the distinction between perception and non-perceptual sensory registration.

I turn now to three points about perceptual representation.

The first is that in the scientific study of perception, there is no basis for claiming that perception involves computation on *non*-contentful *syntactic* items, as is commonly claimed by proponents of the language-of-thought hypothesis. No level of the science postulates a syntax that is separated from its contentful mode of representation. There are studies of the neural underpinnings of perceptual states. And perceptual states themselves are taken to have veridicality conditions, and to represent entities in the environment in structured ways. But there is no intermediate, purely syntactic level of states or processing. No laws are specified in terms of such states or processes. The idea that perceptual psychology *must be* committed to such a purely syntactic level derives from ideology that has no sound scientific or philosophical basis.

It is very important in understanding the language of thought hypothesis, especially as applied to perception, that one separate claims. The claim of the hypothesis that I have been doubting is that perceptual processing involves syntactic items that are individuated independently of representational content. A pair of claims associated with the hypothesis that are *not* committed to this claim are these:

(1) that representational psychological states have structured contents, which, since the contents type-individuate the state kinds, entails that psychological state-kinds are structured;

and

(2) that representational psychological states enter into normal causal relations and can be understood computationally and in terms of their structures<sup>18</sup>

<sup>18</sup> I take this point to entail that transformations in the perceptual system are computational, not just that there are computational descriptions of the transactions. It is common to specify that the causal relations be “local”. This specification has not been clarified for psychological causation. Locality is usually understood in terms of physical proximity. Scientific specifications of psychological causation do not—at least currently—specify place, certainly not with the specificity that would be

I will call the conjunction of (1) and (2) ‘LOT-’.

I accept these two claims. I think that they are supported by the science as well as by common sense. It is, however, misleading to regard these claims as being distinctively about an internal language. They are simply about representational contents, or kinds of states whose individuation involves structured representational contents. There is nothing distinctively linguistic about the phenomena specified by (1) and (2). The arbitrariness of the connection between symbols and their meaning and syntax is, I think, fundamental to being a language.

The claim that I doubt is the following:

(3) that elements in the structure of perceptual states are individuated purely syntactically – in a language-like way that is nevertheless independent of any representational content.<sup>19</sup>

The conjunction of (1), (2), and (3) is the usual language of thought hypothesis, LOT, as applied to perceptual states. I do not accept (3). There is no scientific ground, or I think good philosophical ground, to believe it.

Philosophical grounds have been adduced to support (3). The philosophical issues are extremely complex. I do not hope to do justice to them. But I will indicate briefly why I do not find the philosophical grounds persuasive. (3) is commonly introduced to explain how psychological processes could be realized in a being whose operations

relevant to meet the requirement. I think that one should understand this condition as merely requiring that the psychological causation not invoke anything odd, such as action at a distance.

<sup>19</sup> I take ‘syntactically’ here not to be deflatable to mean simply non-semantically. There are certainly physical structures that underlie perceptual states. But those structures need not have a language-like syntax. I intend a purely syntactic structure to be neither merely a structure described in physical terms nor a structure abstracted from a contentful structure. It must be describable and explainable in recognizeably syntactic, grammatical, language-like terms; and it must ground significant explanation under such description. Syntactical items would be participants in an intermediate level of processing. I doubt that there are purely syntactic structures in perceptual systems, in that sense. I shall discuss the “grammar” of the representational contents of perceptual states shortly. It is an aspect of the representational content, abstracted from the full representational content--not a content-independent kind. I doubt that the physical underpinnings of perceptual psychology have a grammar.

are “sensitive” to purely physical properties. The idea is that the syntactic items are instantiated physically in something like the way word shapes are. Thus (3) is invoked as part of an attempted partial solution to the mind-body problem.

From a mathematical point of view, computability is independent of whether the structured items on which the computations operate are individuated purely syntactically or are (on the contrary) individuated in ways that involve content. Nothing in computability *per se* favors LOT. So nothing in understanding psychological transactions as *computational*, where those psychological transactions are contentful, supports LOT.<sup>20</sup>

As noted, invoking LOT, including (3), is commonly part of trying to solve causal aspects of the mind-body problem. I think that trying to do philosophy of psychology while also trying to solve the mind-body problem is a treacherous enterprise. The dangers have not been well avoided in philosophical discussion. Highly speculative views have been mixed with more concretely scientific views. The speculativeness of the former has commonly been obscured by bluster or ideology. (3) is completely ungrounded in the science of perceptual psychology.

I believe that there is nothing philosophically problematic about computations that operate on the form of psychological states, where the form is constitutively an aspect of representational content. I think that one should resolutely concentrate on how computations are actually referred to in the science. They are operations among the contents of states, or among contentful states. The causation is referred to, in computational explanations in perceptual psychology, as psychological causation, or physical-psychological causation, or psychological-physical causation. Attempts to reinterpret these accounts in terms of sensitivity *purely* to physical properties or to constitutively non-contentful syntactic properties are not supported by the science. Nothing has shown the science to be deficient, or in need of reinterpretation.<sup>21</sup>

<sup>20</sup> Michael Rescorla has independently done work, so far unpublished, making similar points. He establishes in some detail that computation is mathematically definable independently of assumptions about content-independent syntax in ‘Is Computation Formal?’, forthcoming.

<sup>21</sup> Although LOT, and other physicalistically motivated attempts to explain mind-body causation, aimed to illuminate psychological causation, I believe that

Neither dualists nor physicalists have explained, in acceptably clear terms, what the problem of understanding psycho-physical interaction is – let alone what a good solution is. Descartes posed the problem in terms of how substances that are *so* different can interact. This *so different* idea is both vague and dependent on an old-fashioned conception of self-sufficient substance that no one holds (or *should hold*) nowadays. Descartes's opponents have not, in my judgment, clarified what remains from Descartes's worries that we need to worry about.<sup>22</sup>

I do think that there is ground for dissatisfaction with our philosophical understanding of mind-body causation. We do not understand the relations between mental events/properties and physical events/properties. We do not understand causation itself very well. There is plenty to clarify and understand better. But I think that the problem is not that of explaining how a seeming impossibility is possible. The science has shown its possibility by explaining its actuality in systematic law-like ways.<sup>23</sup>

over the last three to four decades they have been notably unsuccessful. They often covertly entail epiphenomenalism. What is scientifically clear is that there is causation that involves contentful states with contentful structures. Such states are individuated and explained partly in terms of their roles in law-like causal patterns. The psychological states have causal powers. These powers produce patterns of computational transitions in which contentful states participate. The causation is “mechanical” in the sense that much of it is quasi-deterministic and computational. Nothing odd about the causation has emerged in the science. Psychological causation itself is not odd, except from reduction-must-be-found points of view. For further discussion, see ‘Mind-Body Causation and Explanatory Practice’ and ‘Postscript to “Mind-Body Causation and Explanatory Practice”’, both reprinted in *Foundations of Mind* (Oxford: Clarendon Press, 2007).

<sup>22</sup> ‘Mind-Body Causation and Explanatory Practice’, *op. cit.* and ‘Postscript to “Mind-Body Causation and Explanatory Practice”’, *op. cit.*.

<sup>23</sup> Contrary to Descartes, it is virtually certain that psychological causation could not go on without underlying physical processes. There is certainly an unbroken chain of purely physical causation that occurs whenever the psychological, or mixed, causation occurs. The latter depends on the former. Whether the physical chain will be seen scientifically and explanatorily as itself organized syntactically in a way that nicely mirrors the structures in the psychological causation is an open question. I bet negatively, even outside the domain of perceptual psychology. Currently, there is no empirical evidence for such a view. I bet that the interrelations between the psychological and the physical are too complex and variable to allow such a neat solution. I think that a content-independent syntax—particularly in perceptual psychology--will never be grounded empirically. Urgent attempts to

What we know is that the psychological laws or law-like patterns involve transformations among *contentful* perceptual states. The formulation of principles stating the laws in terms of states with the representational content (veridicality conditions) is *primitive*. It is not a commentary on a *further*, primitive non-representational syntactical structure. The vehicles of representational content, the perceptual states, are not individuated separately, as a word shape might be individuated independently of its meaning. There is a story to be told about how the perceptual level of specification relates to neural levels of specification. But there is no empirical reason to believe that any states in perceptual systems have a syntactic structure that is determinable independently of the representational contents of the perceptual states.

The representational contents of perceptual states, hence the natures of the states, are constitutively dependent on relations to environmental entities. This point is evident in characterizations of perceptual states *in the science*. The point is an expression of anti-individualism in the science.<sup>24</sup> The science individuates perceptual states as perceptions as of instances of shape, color, motion, body, and so on. No purely syntactical state is individuated in these ways. The science is a science of laws governing formation of perceptual states.

The second general point about perceptual representation is that the principles governing the transformations among perceptual states in perceptual systems are in no sense, not even implicitly, *represented* in the perceptual system. More accurately, the principles are not the contents of any states in the system. Such principles are not “looked up”, “applied”, “accessed” in the system, even “implicitly”. Thinking of the principles in that way would be almost as bad as thinking of the planetary system as applying principles governing its motion. The difference is just that the planetary system contains no representation-

answer Descartes with a *specific* picture have been at best interesting and speculative, but more commonly unscientifically dogmatic. We know that there are dependency relations between the types of causation. Maybe things will work out so that the relations are as simple as LOT proposes. But if they do not, we will be forced toward richer forms of understanding, which transcend the simple pictures that have dominated philosophical discussion since Descartes.

<sup>24</sup> For discussion of anti-individualism, see *Origins of Objectivity*, *op. cit.*, chapter 3.

al states.

Psychological principles that govern transformations in perceptual systems describe real psychological laws and law-like patterns instantiated in the transformations. The information registrations and perceptual states that undergo the transformations are present in the system. The perceptual states are often accessible to consciousness of perceivers, and are certainly attributable to perceivers. That is, not only does the sequence of causation that begins with sensory registration produce a perception of something as of a moving body. But the relevant perception is also not a sub-individual state. The state that is produced *is* the individual perceiver's perceiving or misperceiving something as of a moving body. The states are individuated primarily by their representational contents and thus have those contents. But the contents of the *principles* governing *transformations* that lead to the formation of perceptual states are not the contents of any state in the perceptual system, or of any state of the perceiver, even unconsciously.

A reason not to take principles formulating laws of transformation to be the representational content of even implicit unconscious perceptual states – is that many perceivers lack the resources to specify perceptual states. Perceptual states do not specify perceptual states. They do not specify mathematical structures. Rather the systems contain perceptual states that operate in mathematical patterns specified by psychologists.

The principles governing transformations in perceptual systems are not analogous to rules that are specified, or specifiable, in the proof systems of logicians. They are not even unconscious analogs of such rules. The psychological theories of perception can be modeled on computers. The theories are computable. Perceptual systems are approximately deterministic. They operate computationally. The transitions that occur within perceptual systems are themselves

<sup>25</sup> I emphasize this point because certain misinformed glosses on perceptual psychology take it to be purely a science of sub-individual psychological states, not a science of individual (including human) perception. To be sure, there are many sub-individual states in any perceptual system. The transitions in the system are nearly all irretrievably unconscious and sub-individual. But much of the point of the science is to include individuals' perceptual states, including phenomenally conscious and consciously accessible perceptual states, in a system in which formation events and law-like patterns of formation are explained in mathematically rigorous ways.



computational. The information registrations that are inputs into the system together with the laws of transformation determine, *modulo* interferences, perceptual outputs, at least at early levels of perceptual processing. None of these points, however, supports the view that the principles of transformation are the contents of any psychological state of the perceiver, or any sub-individual state in the perceptual system.

The third general point about perceptual representation is that it *does* have a definite representational structure, analogous to a grammar or a logical form. Recall the first point – that there is no empirical reason to maintain that there is a *pure* syntactical level of states or processing. The point was that there is no level that is independent of the representational contents of the states. The basic states in perceptual systems, perceptual states, are constitutively representational. It does not follow from this first point that the representational content of perceptual states lacks a representational structure or form. It is just that the form (the “syntactical” form) is an aspect of the representational contents of the states.

I believe that, at the most generic level, the representational structure of perceptual states can be determined from considerations about perceptual function and perceptual veridicality conditions. All perception must have representational content that contains both general attributive elements and singular referential elements.<sup>26</sup>

I begin with the general elements – what I call ‘*perceptual attributives*’. Perceptual states must represent what is perceived as being a certain way. There is no such thing as perceiving anything *neat*. Perception types, characterizes, groups, attributes. I take it that perceptual contents *indicate* properties, relations, or kinds. But they do so only while attributing those properties, relations, or kinds to purportedly perceived particulars. One cannot *just* perceive a particular, or even a purported particular. One must perceive it as being a given color, or shape, or as being at a certain distance, or as being a moving body. Perception attributes these indicated attributes (property, relation, or kind) to purportedly perceived particulars. Even the attributes that are perceptually indicated must be represented in a

<sup>26</sup> Perhaps there are pluralized referential applications in perception, with something like the form of ‘those’. I shall, however, mainly write of singular applications. There are always many such referential applications in any given perceptual content.

certain general way. One can perceive something as circular utilizing either the head-on way of perceiving an instance of circularity or some way that involves a perspective on the circle from an angle. Such ways of indicating attributes and attributing these attributes are general forms of representational content – perceptual attributives.

The general elements in perception – perceptual attributives – are general in two ways relevant to our discussion.<sup>27</sup> First, they are general in marking, or in helping to type individuate, general abilities – abilities that are not constitutively dependent for their identities on any particular occasion of use. All perceptual attributives that are general in this first way are also general in a second way. They are general in that their representational function is to be *veridical of* one or more entities. They have an attributional form and role.

Thus perceptual attributives have a psychological role in individuating a certain aspect of a perceptual capacity or perceptual state. Different perceptual attributives may indicate the same attributes and be veridical of the same particulars. They differ in that they help mark different kinds of perceptual states or capacities. The capacity or state marked by an attributive that indicates circularity and attributes circularity to particulars from a head-on perspective is a different capacity or state from one that indicates and attributes circularity from a given angle. It is an achievement of a perceptual system to track an instance of circularity as such, even though the instance is presented in such perceptually different ways.

I call the singular elements in perception '*singular (or plural) perceptual applications*'. These elements are, like the general elements,

<sup>27</sup> There is an element in the representational structure of perceptual representational contents – which I shall not focus on here--that is general in the first of these senses, but not the second. These are schematic elements analogous to the demonstrative 'this', considered as unused. Such elements mark a general capacity to apply perceptual attributes to particulars. Although such elements mark general capacities, they are not veridical of anything. They are syntactically singular, even though they can apply to any number of particulars (they are semantically general) and even though they mark a general ability (they are ability general). The first type of generality mentioned in the text is what I call '*ability generality*'. The second is a species of what I call '*syntactic generality*'. I believe that perceptual attributives are also semantically general, but I shall not discuss this point here. For more concentrated discussion of these different types of generality, see my 'Five Theses on *De Re* States and Attitudes' in J. Almog and P. Leonardi eds., *The Philosophy of David Kaplan* (Oxford: Oxford University Press, 2009).

discernible by reflection on the function and veridicality conditions of perception. Perception functions to be of particulars, not merely repeatables or types. The particulars that it succeeds in being of are *concrete*. They are non-repeatable entities in time that cause the occurrence of successful perceptual states. Perception also functions to pick out, identificationally single out, particulars in a direct, not purely attributive manner. It functions to refer to them.

Perception sometimes fails its function. In such cases, it still *functions* to refer to particulars: it is *as of* particulars that it functions identificationally to single out. Perception functions to refer to particulars as being of general types.

Perception's functioning to single out particulars figures in the veridicality conditions of a perceptual state. Whether the state is accurate or not hinges on whether it succeeds in singling out relevant particulars and, if it does, whether it attributes to those particulars properties, relations, or kinds that the particulars in fact have or instantiate. Thus the representational content of perception constitutively has singular (or plural) referring elements. These elements are singular both in that they mark occurrent exercises of abilities on particular occasions and in that their representational function is to refer to, or single out, entities.

Perception singles out particulars in context-dependent ways – ways that depend on a causal context, as well as on perceptual attribution. So the singular elements are context-bound. Singular perceptual applications are the perceptual analogs of occurrent *uses*, on particular occasions, of singular demonstratives in language.

The singular elements must always operate together with the general elements. They are, in fact, referential applications of the general elements to purported particulars. They focus the general elements' attributive powers on concrete particulars in a given causal context. When the singular elements are representationally successful – when they fulfill their representational function – they not only pick out particulars; they also apply the general elements to particulars that the general elements function to be veridical of. Here, as always, psychological form is to be explained in terms of representational function. Representational function is grounded in representational competencies – the basic competencies involved in perception.

Singular elements are not immune to failure of reference. General attributives are not immune to failure of indication, or failure of attribution. In conjunction, they constitute the veridicality conditions

of a perceptual state. The generic form of a perceptual state is, crudely,  $\text{that}_i E$ , where the subscript marks an occurrent application of the demonstrative-like, syntactically singular form ‘that’.<sup>28</sup>

It is not a requirement on the existence or referential success of these singular applications that they (or perceptual contents containing them) be conscious. We know that insects have perception, with a large array of perceptual constancies. We do not know whether insects are conscious. In fact, there is evidence that some color constancies – hence perceptions – in bumblebees occur at the retinal level, with nearly no processing. Such constancies almost surely occur before consciousness could occur, even if the bees *are* conscious.<sup>29</sup>

There are also, in humans and other animals, certain states formed in the first microseconds of visual processing that seem to involve perceptual constancies, hence perception by individuals, but seem not to be conscious. Again, it is likely that these constancies are formed before any kind of consciousness can occur. Such states are, at the very least, not consciously *accessible*. The individuals are oblivious to what they perceive.<sup>30</sup>

Moreover, we know that blindsight patients perceive entities, again showing an array of perceptual constancies. It is very likely that

<sup>28</sup> Some of the crudities: In addition to the schematic element in the general form ‘that’ (see note 27), I have not discussed differences between perceptions of instances of properties and relations, and perceptions of entities, such as bodies or events, that *have* the properties and figure in the relations. I have not discussed the embedding of perceptual content in ego-centrally anchored frameworks. I have not discussed the fact that perceptions involve multiple references to multiple particulars (purportedly instantiating multiple attributes) at different levels of abstraction—the fact that perception is saturated with multiple references. Since my main concern is to sketch the rudiments of the structure of perceptual content, I do not go into these matters here.

<sup>29</sup> Adrian G. Dyer, ‘Bumblebees Directly Perceive Variations in the Spectral Quality of Illumination’, *Journal of Comparative Physiology A* 192 (2006), 333-338.

<sup>30</sup> Steven J. Luck, Edward K. Vogel, and Kimron L. Shapiro, ‘Word Meanings Can Be Accessed But Not Reported During the Attentional Blink’, *Nature* 393 (1996), 616-618; Stanislas Dehaene, Lionel Naccache, Guryan Le Clec’H, Etienne Koechlin, Michael Mueller, Ghislaine Behaene-Lambertz, Pierre-Francois van de Moortele, and Denis Le Bihan, ‘Imaging Unconscious Semantic Priming’, *Nature* 395 (1998), 597-600; Rene Marois, Do-Joon Yi, and Marvin M. Chun, ‘The Neural Fate of Consciously Perceived and Missed Events in the Attentional Blink’, *Neuron* 41 (2004), 465-472.

the relevant perceptions by blindsight patients are not phenomenally conscious. Blindsight is just one of many types of dissociation in which unconscious perception occurs. Prosopagnosia and extinction-neglect syndromes are others.<sup>31</sup>

There is very strong evidence that perception by individuals of particulars as having specific attributes – perception exhibiting numerous perceptual constancies – occurs, even though the relevant perceptual states are not phenomenally conscious, and even though the individual is not conscious of the perceptions, in the perceptions, or of the entities being perceived. It is simply a mistake to assume that all genuine perception is conscious.<sup>32</sup>

A focus on conscious perception can be legitimately motivated in

<sup>31</sup> I discuss these matters in greater depth in *Origins of Objectivity*, *op. cit.*, esp. 374-376. For a sampling of relevant psychological literature, see L. Weiskrantz, *Blindsight* (New York: Oxford University Press, 1986); R.W. Kentridge, C.A. Heywood, L. Weiskrantz, 'Spatial Attention Speeds Discrimination Without Awareness in Blindsight', *Neuropsychologia* 42 (2004), 831-835; James Danckert and Yves Rossetti, 'Blindsight in Action: What Can the Different Sub-types of Blindsight Tell Us about the Control of Visually Guided Actions?', *Neuroscience & Biobehavioral Reviews*, 29 (2005), 1035-1046; Daniel L. Schacter, Mary Pat McAndrews, and Morris Moscovitch, 'Access to Consciousness: Dissociations between Implicit and Explicit Knowledge in Neuropsychological Syndromes', in *Thought Without Language*, L. Weiskrantz ed. (Oxford, Clarendon Press, 1989); Martha J. Farah, 'Visual Perception and Visual Analysis After Brain Damage: A Tutorial Overview', in C. Umiltà and M. Moscovitch eds., *Attention and Performance XV: Conscious and Nonconscious Information Processing* (Cambridge, Mass.; MIT Press, 1995), 37-75, also in N. Block, O. Flanagan, and G. Güzeldere eds., *The Nature of Consciousness*, (Cambridge, Mass.; MIT Press, 1998). See also Bruce T. Volpe, Joseph E. Ledoux, and Michael S. Gazzaniga, 'Visual Processing of Visual Stimuli in an 'Extinguished' Field', *Nature* 282 (1979), 722-724; M. Verfaellie, W.P. Milberg, R. McGlinchey-Berroth, L. Grande, and M. D'Esposito, 'Comparison of Cross-field Matching and Forced Choice Identification in Hemispatial Neglect', *Neuropsychology* 9 (1995), 427-434; James P. Morris, Kevin A. Pelphrey, and Gregory McCarthy, 'Face Processing Without Awareness in the Right Fusiform Gyrus', *Neuropsychologia* 45 (2007), 3087-3091.

<sup>32</sup> To reverse the point, not all consciousness involves representation in my non-deflated sense. Awareness of the felt quality of pain (as distinguished from proprioceptive locating of pain) does not require representational content, perceptual constancies, perceptual objectification, or any of the other marks of genuine representation. There may be organisms that feel pain—and hence are conscious—but lack any representational capacities (again, in the non-deflated sense of 'representational').

various ways – interest in the phenomenology of perception, interest in special powers of the consciousness in perception, interest in individual responsibility for perceptual information (which plausibly requires some sort of consciousness). But the assumption, often made in philosophy, that perception must be conscious to be “genuine” is empirically mistaken. Such an assumption tends to hold back progress in philosophical understanding of both the psychology of perception and the epistemology of perceptual belief.

There are difficult questions about the role of consciousness in perception. This issue is not yet well understood scientifically. I conjecture that phenomenally conscious elements are an aspect of the mode of presentation, the representational content. Conscious aspects of perceptual states render those states psychologically distinct kinds of states from unconscious counterpart states that attribute the same attributes to the same particulars. Since perceptual states are individuated in terms of their perceptual mode (vision, hearing, touch, proprioception) and their representational content, it is plausible that phenomenal consciousness is, in these cases, an aspect of the representational content of perceptual states. For example, a conscious perception of something as cubic is, I think, commonly a different perceptual state, with a different representational content, from an unconscious perception of the same thing as cubic. These are different ways of perceiving something as cubic – ways that bear on the individuation of the state-kind. In such cases, it is misleading to say that a given psychological state came to consciousness, or went out of consciousness. Only a generic kind is shared.<sup>33</sup>

It is also not a requirement on perceptual reference to particulars that the perception involve attention. The previous point about the independence of perception from consciousness is relevant here. But attention itself can be unconscious. Neither conscious nor unconscious attention is necessary for veridical perception of particulars as having specific attributes. Individuals perceive instances of attributes and rely on perceptual constancies even when they are not attending to the perceived particulars. Although certain types of multiple object

<sup>33</sup> My remarks here barely touch a complex set of issues. I think that differences in *degrees* of consciousness are not *in general* differences in modes of presentation.

tracking require attention,<sup>34</sup> much perception of particulars as having or instantiating attributes occurs even though the perceived entities are not attended to.

There is empirical reason to believe that perceptual representational content *per se* does not figure in propositional inference. There is therefore empirical reason to believe that the perceptual representational content is not propositional. I think that attribution of propositional form to perception plays no role in the science. There is strong empirical reason to hold that perceptual content is not propositional. Perhaps any representational structure can be mapped onto a propositional structure. But not all psychological representational structures are propositional. Representational structure marks representational kinds. There is reason to believe that a representational kind is propositional only if propositional structure – particularly propositional inferential structure – grounds explanation. I think that such structure plays no role in grounding explanation of the psychological behavior of perceptual states.

Further, conditionals, disjunctions, quantifiers seem clearly missing from the representational content of perceptual states. I believe that negation is also absent. Cases of perception of absence, or purported cases of negations in perception, are properly accounted for in other ways than by attributing propositional negation to the representational content of the perception. For example, one might be inclined to say, ‘the spider sees that a given navigational line is not a route to the prey’. The spider perceives some break in the route, and has an unfulfilled perceptual anticipation as of a continuous route (or some unfulfilled search for a route that meets some template of continuity). The failure of match between the perception of a spatial break in the route with the perceptual anticipation of continuity (or the search for continuity) can ground explanation without taking negation to be part of the perceptual content. Perception always yields a positive, committal representation of the environment. Rejection, disappointment, lack of fulfillment pertain to modes of responses to perceptual presentations, not to the representational

<sup>34</sup> Z. W. Pylyshyn, ‘Visual Indexes, Preconceptual Objects, and Situated Vision’ *Cognition*, 80 (2001), 127-158.

content of perceptual states.<sup>35</sup>

In fact, the explanation of the formation of perceptual states in perceptual psychology does not appeal to the formation of states with propositional content. Nor is there any need in the theory to invoke propositional structure in perceptual content. The structure of the perception contains occurrent, context-bound applicational elements (singular or plural) and attributional elements that attribute repeatable attributes to particulars. No further structure is appealed to or needed in scientific explanations of perceptions. There is therefore empirical reason to believe that perceptual representational content is not propositional.

In ordinary language, we use locutions like ‘She has a perception that there is a red car directly ahead’. Such locutions can describe real states. Such locutions apply to hybrids of natural psychological kinds. There are two points to be made here.

One is that attributives like car are not in the repertoire of basic perceptual systems *per se*. Basic perceptual systems are language independent. They are systems of perceptual competencies that are relatively self-contained and have substantial innate aspects. Car enters a psychology through language. What we correctly call perceiving something as a car is what is known in the science as an instance of *high-level perception*. High-level perception occurs through the integration of post-perceptual capacities and perceptual capacities, strictly so-called. The post-perceptual attributive car is applied through its association with entities perceived as bodies with characteristic shapes, colors, and perhaps motion. The repertoires of perceptual systems, and of *perception* by individuals scientifically so-called, include attributives for shape, color, motion, body, spatial relations, temporal relations, and various other attributes, possibly including functional attributes such as danger, food, or fear. But they do not include attributives for cultural attributes like being a car. I know of no a priori principle here. The point is empirical. But it is well-established.

<sup>35</sup> Given that there is empirical reason to reject the view that perceptual content is, at the fundamental level, propositional, there is reason to deny that any “conjunctions” in perception are propositional, logical connectives. Commonly, “conjunctions” in perception amount to plural attributions to a single particular.

<sup>36</sup> Shimon Ullman, *High Level Vision: Object Recognition and Visual Cognition* (Cambridge, Mass.: MIT Press, 1996).



The second way in which the has-a-perception-that locution describes a hybrid is that visual perception is not propositional. The ordinary perception-that locutions can (where they impute only perceptual attributives) be understood as loose descriptions of instances of the psychological kind *visual perception* – descriptions that do not purport to convey the form or structure of the perceptual content. But more commonly they are best understood as form-fitting descriptions of apropositional attitude whose conceptual attributives (including but not limited to ones for body, shape, and so on) are connected with a perceptual basis, strictly so-called, for applying these conceptual attributives. Ordinary language is not fitted to science. But it is not completely out of touch with the natural psychological kinds that science uncovers.

Some have thought that whether perception is propositional depends on whether what is represented is an object or a state of affairs. I believe that this view has things backwards. The issue is over the organization, structure, form of the representational state. A state of affairs can be represented either with a singular representational structure or with a propositional structure. Similarly, a relation between different particulars can be represented in a nominal or propositional way. The question at issue concerns the *way* – the structure of the mode of presentation, of the representational content. That structure hinges on the representational function and psychological capacities involved in perception. The function is referential and identificational: the structure is governed by demonstrative-like applications of singular (or plural) perceptual attributives.

Broadly speaking, I believe that the form of perceptual content is more like that of a noun phrase than that of a sentence. Of course, basic perceptual contents are not conventional, as the relation between noun phrases and their contents is. They cannot be manipulated at will by the perceiver, as noun phrases can be. And ground-level perception represents more or less continuously and concretely, unlike noun phrases, but like at least some maps. But the function of perceptual contents, like singular noun phrases, is to single out particulars, by causal-contextual connection to them and by determining their attributes. All other representational functions of perception are subordinate to that one.

This function shows up in the representational ways in which per-

ception can go wrong. Perceptual states can go wrong if they fail to refer to a concrete particular, a particular in time that causes the perception in an appropriate way. Failure of perceptual reference undermines success for the other main representational aspect of the perceptual state – the attributional aspect. All the attributions that function to be accurate of a purportedly perceived particular fail to fulfill their function, if no concrete particular is referred to – if no concrete particular is perceived. Most of the attributions can go wrong even if a particular *is* perceived.<sup>37</sup> They go wrong if they are not accurate of – do not accurately characterize – the perceived particular. Those are all the ways that perceptions can go representationally wrong. In these respects, the form and representational evaluation of perception are like the form and representational evaluation of applied demonstrative-governed noun phrases. Representational failure is a failure to fulfill representational function. The representational form of perception reflects its representational function. Again, form follows function.<sup>38</sup>

<sup>37</sup> I say ‘most’ because there is, I think, always an attributive applied in any perceptual reference or perceptual indication that must be accurate of a perceived particular, if the particular is to be perceived at all. For example, to perceive a body or shape at all, there must be something in the distal environment that appropriately causes the perception and that meets certain topological conditions of connectedness attributed in perception. If the percept is caused only by a scatter of light particles that form no connected shape, then no body or shape is perceived. There is only a referential (or indicational) illusion of body or shape. With respect to most perceptual attributives that are applied in perception, perceptual reference does not depend essentially on the attributive’s being accurate of the particular that is perceived. One can perceive the entity but misperceive it as having most of the attributes that one perceives it as having.

<sup>38</sup> This paragraph indicates that the pattern of possible failures of representational function is consistent with and suggestive of the representational function and representational form of perception. It is *not* meant to produce an argument that perceptual states have a nominal rather than propositional form. The key point in the nominal/propositional distinction is that in perception all attributive functions are subordinate to referential function, whereas in propositional states, an attributive function is free from referential function. That point requires more explanation. There are delicate issues here in explaining what functions and competencies underlie the difference between the nominal (noun-phrase-like) form of the representational content of perceptual states and the propositional form of *simple* perceptual belief: perceptual beliefs with no logical constants. The representational function of the main predicative attributive in the propositional contents of simple perceptual beliefs is *not* subordinate to reference. It functions purely attributively

The singular applications inevitably apply general attributives. In fact, the singular applications are applications of attributives to the particulars. Particulars include instances of properties and relations, as well as instances of kinds (such as *body*). Attributives function to indicate attributes (specific color, shape, motion, bodihood, and so on) and to be veridical of, and attributed to, the particulars that perception singles out.

Obviously, many propositional contents contain singular applications; and all propositional contents contain attributives. How are we to ground the difference between the structure of perceptual representational content and propositional structure? I think that part of the answer is that every perceptual attributive accompanies and is applied by a contextual, demonstrative-like (context-bound) singular (or plural) perceptual application.<sup>39</sup> Again, the representational function of perception is to single out and characterize particulars. All perceptual attributives function in (and are subordinated to) context-bound, identificational reference to particulars. The fact that attributive competencies are *never* exercised separately from referential, identificational, singular or plural applications in perception signals the non-propositional status of perception.

In propositional structures there is *always* a predicative element with attributive potential that does not function to accompany and be applied by singular reference. There is *always* an attributive that functions predicatively outside the scope of any identificational demonstrative-like element. The main predicate (the analog of the main verb phrase in sentences) is never within the scope of an identificational element. For example, in the propositional content That<sub>1</sub> body is spherical (applied in a context), the attributive is spherical does not function to help single out anything – it is not applied by a singular element – whereas the attributive body does function to help single out something; it functions to single out a body. Is spherical is not

and purely predicatively. This difference is reflected in the fact that the main predicative attributive is not within the scope of a referential application. Explaining this difference in detail (in particular, in terms of underlying representational capacities) is, I think, a complex and profound matter. I flag the issue here and again one paragraph down in the text. I discuss it further in the next lecture, but confront it head-on only in yet further work.

<sup>39</sup> I explicate these notions, under the headings ‘guidance’ and ‘accompaniment’ in ‘Five Theses on *De Re* States and Attitudes’, *op. cit.*. I now prefer my present terminology.

applied by the singular element. It is not within its scope. All propositional structures have main predicates. All attributives in propositional structures – what I call ‘*conceptual attributives*’ – are capable of occurring as main predicates. Moreover, the attributive body can in other propositional contents function purely to make an attribution while not being applied by context-bound referential elements – as in Every planet and moon in the universe is a body. These are all cases of *pure attribution* with the attributive body.

In yet further structures, the attributive body occurs predicatively, but does not make an attribution at all. I have in mind the propositional contents It is not the case that that<sub>1</sub> form is a body and If any form in the room is a body, it has mass. These are cases of *pure predication*. Here the predicative use and representational function of the attributive are independent of attribution, let alone attribution that is applied in referential identification.

Perception has no purely attributive or purely predicational uses. It conspicuously lacks any representational function that is independent of referential identification. The representational function of perception is identification. The form of perceptual representation – and of its representational content – reflects its representational function. All perceptual attributives function within identifications. All perceptual attributives occur within the scope of demonstrative-like referential, identificational applications of those attributives.<sup>40</sup>

I have distinguished perception and propositional states both structurally and functionally. Structure and function must ultimately be explained in terms of competence and exercise of competence. So we need a deeper explanation, in terms of differences in competencies and their uses, that grounds the differences in structure and function.

Most philosophers in the twentieth century took the problem of explaining how objective singular reference emerges in propositional thought to be the problem of explaining how conceptual capacities,

<sup>40</sup> As I emphasize in the second lecture, ‘Steps Toward Origins of Propositional Thought’, one should not conflate this point with the distinction between attributives that are *essential* to perceptual reference and attributives that are not. Few perceptual attributives are essential to perceptual reference (see note 37). But all perceptual attributives function to be applied by referential, singular or plural perceptual applications. They all have no further representational function than to be applied attributively to referred-to (perceived) entities. Their attributive function is subordinate to and serves the basic identificational, referential function of perception.

even linguistic capacities, make possible singular reference. But objective singular reference to entities in the environment is already present in autonomous perception, perception in animals that surely lack propositional abilities.

Given that singular objective reference is present in pre-propositional animals, like insects, the real problem is to explain the liberation of attributives from functioning entirely to be applied by singular context-bound identificational reference. The problem is that of explaining the competencies that ground the separation of pure attribution and pure predication from perceptual attribution. Pure attribution and pure predication are attribution and predication that do not function attributively within the scope of context-dependent, *de re*, identificational reference.

Conceptual attributives, or predicates – as I use these terms – constitutively contribute to propositional contents. They are constitutively capable of predication that does not function in identification – pure predication.<sup>41</sup> To have conceptual attributives, an individual must be able to use attributives in ways other than in identifying particulars. I believe that this ability is closely associated with the emergence of propositional inference.

Perception constitutes the origin of representational mind. As a matter of empirical evolutionary fact, perception seems to show up first among the arthropods. A later, momentous origin is the advent of propositional attitudes and reason. I am interested in the constitutive conditions that go with this later origin. In further work, I reflect

<sup>41</sup> More generally, *concepts*—in my usage—are constant, freely repeatable elements in propositional representational contents. See ‘Five Theses on De Re States and Attitudes’, *op. cit.*. Conceptual attributives—predicates—comprise one type of concept. They are the analogs, in thought contents, of linguistic predicates. Other types of concepts are functional notions (like the successor of), individual constants (2), logical constants (or), and constant schematic demonstrative elements (I or that—as distinguished from their occurrent applications in a context). Context-bound (occurrence-based) singular or plural applications are non-conceptual elements in propositional contents. Analogs of conceptual attributives and context-bound singular (or plural) applications occur in perception. There are perceptual attributives and perceptual occurrence-bound singular (or plural) applications. Attributives in *propositional* structures differ from attributives in perceptual structures in being associated with a wider variety of functions and competencies. Context-bound applications in propositional thought differ from their perceptual counterparts in what sorts of attributives they apply. In the next lecture, I shall take these points somewhat further.

on the transition from perceptual powers, and other identificational powers, to propositional powers – including powers of propositional inference and reason.