Perception: Where Mind Begins

TYLER BURGE

Abstract
What are the earliest beings that have minds in evolutionary order? Two marks of mind are consciousness and representation. I focus on representation. I distinguish a psychologically distinctive notion of representation from a family of notions, often called 'representation', that invoke information, causation, and/or function. The psychologically distinctive notion implies that a representational state has veridicality conditions as an aspect of its nature. Perception is the most primitive type of representational state. It is a natural psychological kind, recognized in a mature science: perceptual psychology. This kind involves a type of objectification, and is marked by perceptual constancies. The simplest animals known to exhibit perceptual constancies, perception, and representation in a distinctively psychological sense, are certain arthropods. Representational mind, or representational psychology, begins in the arthropods. We lack scientific knowledge about the beginnings of consciousness. Consciousness is neither necessary nor sufficient for perception. I conclude by reflecting on the kinds *mind* and *psychology*.

Where does mind begin? This seems like a natural question. Rocks and fires, floating in empty space, are overwhelmingly the dominant large citizens of the universe. Most of us are confident that rocks and fires do not have minds. *We humans* have minds. Do any other terrestrial beings have minds? If so, which ones?

There is a philosophical question about how each of us knows that others have minds. That is not my question. I refuse to worry about whether other people have minds, at least not in this talk. I will not try to answer sceptics about mind in general or about how knowledge of minds is possible. I am more interested here in what we know than in sceptical questions about whether we know, or how we know.

I do have some sympathy with scepticism about whether some *particular* humans – certain politicians for example – have minds. But I will not be dealing even with these sceptical questions. I will assume that we do have some knowledge that will help answer our question. My task is to explain that knowledge.

So our question is: Do other types of terrestrial animals have minds? If so, which ones?

Most of us think that apes and dolphins have minds. And cats and dogs. The cats seem willful. The dogs seem to want to be with us. Both have eyes that seem to express mindfulness. And we hear all

doi:10.1017/S003181911400014X

the time about how smart apes and dolphins are. Willfulness, wanting, expressiveness, smarts all seem to be signs of mind.

What about birds, with their bird brains? What of fish with their lifeless eyes? What of snakes with their robotic, mindless-seeming reflexes?

What of spiders, ants, bees? They lack the endearing eyes that cats and dogs have. Like snakes, the bees and the spiders seem to act in 'pre-programmed' ways. Their repertoire of behavior is limited. What they do is relatively easy to predict. So maybe insects lack minds.

But the ants and bees do seem to cooperate in making and doing stuff. They seem to act in each others' interests. Aren't interests things that minds have? Spiders know how to spin webs, and know how to get away when we threaten them. Surely knowledge is something that occurs in minds. Bees communicate with one another. Doesn't communication involve some mental state? So maybe bees and ants do have minds.

What of worms? What of jellyfish? What of amoebae, paramecia, bacteria? They move around purposively and seem to find places where they thrive and avoid places where they do not thrive. They seem to know how to navigate. Purpose and knowing seem to be states of mind.

Well, what of plants?

Aren't we going overboard? But wait. Plants grow in ways that are purposive and beneficial to them. Recently a biologist claimed in *Scientific American* that trees see. Surely seeing requires a mind. Are we being parochial to exclude plants from having minds? Maybe excluding plants is just a prejudice that one day we will look down on. Why shouldn't purposeful growth count as a kind of planning and realization of self-interest? Why shouldn't knowledge and perception be ascribed to plants?

But then, there is a regularity and order in all of nature. Maybe that order is mindful or mindlike. Maybe rocks *like* to move in the ways they do. Maybe fires *enjoy* their dancing. Do rocks and fires have minds after all?

Somewhere in this line of questioning, it is easy to feel that things have gotten out of hand. Some of the questions seem silly. But it is also easy to feel that many of the questions aren't easy. People disagree about how to answer some of them.

Even where there is agreement, it is not easy to say why we agree.

Are the 'right' answers just a matter of cultural habit or choice? Are they relative to what one feels for, or cares about? Are they just a

matter of what stance one decides to take up toward other beings in the universe?

I think not. I think that there are definite answers, based on scientific knowledge.

In asking the questions, one is likely, if one is honest, to feel uneasy. What is one asking exactly? What *is* mind? Is one so sure of what one means? Is it so evident *what mind is*? If the meaning of the question is unclear, how can we even begin to answer it? For all that, there remains a sense that the initial question is a real question, and that it should have a real answer.

Philosophy takes up large, interesting-sounding questions that on reflection need to be clarified before they can be well answered. The larger the question, the more need for clarification. Commonly, philosophy must worry about the meaning of its questions at the same time that it tries to develop answers to those very questions.

That seems like a paradox to some. How can one answer a question that one doesn't understand?

In fact, a lot of our thinking involves terms or concepts that we do not understand very well. We talk about semi-conductors, electric fields, radio waves, lymphomas. We understand these terms well enough to use them. But many of us cannot explain them in any depth. So the first thing to remember here is that it is not so uncommon to ask questions that we don't understand very well. Part of finding the answer is improving our understanding of terms in our questions. Questions about mind are like that.

But don't we need *definitions* to ask a serious question about what individuals have minds? People often ask philosophers 'what is your definition?' for this or that. Here, the question might be, 'what is your definition of "mind"?'.

Philosophers rarely give definitions. They are not shirking their duty. Definitions are not what they are commonly cracked up to be. In fact, we know from some very good work in philosophy by Quine and Putnam that useful definitions are not secure starting places. They can incorporate mistaken theories. Even when they are invoked in science, they are open to questioning and improvement.

We can have security if we *stipulate* meanings. But stipulation yields no new understanding. To be useful, definitions must capture something that we already understand. To be accurate, they must accord with facts. They do not come for free. In fact, they belong, if anywhere, more toward the end of inquiry than at the beginning.

How are we to attack our problem? I will start with some very general remarks about mind. 'Mind' is not a well entrenched scientific term. Ordinary speech allows quite a range of applications. My initial discussion brings out how unsteady ordinary talk is when it comes to deciding whether certain animals have minds. We say that worms know things (that safety is downward into the earth) and want things (to be left alone). But we also see worms as mindless organisms driven by their biological programs.

The tradition of reflecting on mind in philosophy can help. The tradition has focused on two large features, each of which has been taken to be a mark of mind. Some philosophers have taken all minds to involve both features. Others have favored one over the other.

The two large features are *consciousness* and *representation*.

Consciousness needs little introduction. It is what it is like to be awake. If you *are* awake, you probably get it. Notice: no definition.

There are unconscious aspects of mind. In fact, most of what goes on in our minds is unconscious. But the unconscious parts serve the conscious parts. And an individual that is not *capable* of being conscious counts, for many philosophers, as lacking a mind. So for many philosophers being *capable* of consciousness is a mark of having a mind.

I will not be saying much about consciousness, though I will return to it at the end. I won't say much because, although a lot of philosophy and some science have gone into studying consciousness, consciousness is not very well understood. We do not know where, in the evolutionary process, consciousness begins. We can be confident that we are conscious, that apes, monkeys, dolphins, whales, dogs, and cats are conscious. We can be confident because we know that consciousness depends in some way on the brain–for example sensory centers like pain centers. And we know that these animals have brains that are similar enough to ours to infer that they are conscious. They feel pain. We do not know whether insects or worms are conscious. In fact, we do not yet know how to determine scientifically whether they are conscious.

The other large feature of mind is *representation*. I will focus on representation.[1] The term 'representation' is used in several ways. I

[1] Philosophers used to call representation 'intentionality'. This term was closely associated with some bad philosophical theories, and it is easily confused with being intentional in the sense of being on purpose. So I prefer the term 'representation'.

want to clarify a use that centers on a distinctively psychological or mental kind.[2] *Representational* psychological states are those that have veridicality conditions *as an aspect of their natures – as an aspect of fundamental explanation-grounding kinds that they instantiate*.

Let me provide some commentary on the ideas in this point.[3]

First, veridicality conditions. Being veridical is being right about a subject matter. *Veridicality* has two main species – *accuracy* and *truth*. I use the term 'truth' for the veridicality of propositions. Propositions, or propositional representational contents, are simply things that can be true or false and that have the same structure as complete sentences or complete thoughts.[4] I use the term 'accurate' for veridicality that is not propositional. Symbols that are accurate or inaccurate are drawings and maps. Perceptions are examples of psychological states that do not have propositional structure and that are accurate or inaccurate, but not true or false.

A veridicality condition is a condition for being right about a subject matter. A representational content is such a condition. If the condition is fulfilled, the representational content (whether it be propositional or non-propositional) is veridical. If the condition is not fulfilled, the content is not veridical. A propositional representational content is true or false. For example, a belief that dolphins are friendly to humans has a truth condition, a propositional representational content, that is fulfilled if dolphins are in fact friendly to humans, and not fulfilled otherwise. Since representational contents are kinds of psychological states, the associated kind of psychological state – in this case, the belief – is true or false, as well.

[2] Although some philosophers distinguish the kinds *psychology* and *mind*, I begin by taking them to be the same kind. The notion *mind* is more often associated with consciousness. The notion *psychology* (as in having a psychology) is more often associated with representation. Descartes is an early source of this distinction. I shall return to it at the end of the talk. Through most of the talk, however, I will not be drawing any such distinction.

[3] Much of what follows in this talk is elaborated in much greater detail in my *Origins of Objectivity* (Oxford: Clarendon Press, 2010), chapters 8 and 9; and in my 'Origins of Perception', *Disputatio* **4**:29 (2011), 1–38.

[4] The reference here to sentences is just illustrative. Sentences have and exhibit propositional structure. I do not think that propositions are sentences, or that all propositional contents are expressed by sentences. I take the notion of propositional structure to be more basic than the notion of structure of complete sentences.

Similarly, a non-propositional representational content sets a condition for being accurate about a subject matter. If the condition is met, the content and the associated perception are accurate. If the condition is not met, they are not. Representational contents of perceptions are kinds of perceptions. A perception of something as spherical is accurate if that something is spherical. Otherwise, the perception is not accurate.

So much for exposition of the notion of veridicality condition. The rest of my explication of the term 'representational' holds that to be representational a state must have veridicality conditions *as an aspect of its nature – as an aspect of fundamental explanation-grounding kinds that it instantiates*.

Nearly any state can be treated instrumentally, as if it had veridicality conditions. Doing so enables us to treat states as psychological or mental when we do not think that they really are. We can say that a planet wants to get around the sun in the usual way; and its want is fulfilled, made true, if it does so. Bacteria can be treated as more or less accurately representing light or oxygen, and as knowing light or oxygen to be harmful and wanting to get away from them.

Treating bacterial states as if they had veridicality conditions can be useful, but treating them that way does not make it so. In the psychologies of some individuals, having veridicality conditions is an aspect of the *natures* of the states – the kinds of states they are. A scientific understanding of the states themselves invokes veridicality conditions. The veridicality conditions are aspects of the kinds of states that are involved. The laws that the states figure in are specifiable partly in terms of veridicality conditions. Science does not explain bacterial movement in terms of veridicality conditions. Bacteria do not think or perceive. Science does explain the formation of certain psychological states – most impressively perceptual states – in terms of veridicality conditions. Certain animals, including human beings, have veridicality conditions – kinds of representational states – as real aspects of their psychologies.

As noted, it is often convenient, though metaphorical, to invoke veridicality conditions in a description of a phenomenon. Some have taken having veridicality conditions as simply a status accorded those states for the purposes of someone who describes the states. Such a view is sceptical that veridicality conditions, or representation, in the strong sense that I have laid out, are real psychological kinds. Such a view holds that representationality is, in general, a matter of

our stance toward a state, not a real feature of the state.[5] 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. The most developed science of this sort is perceptual psychology. There is empirical reason to take representation, in our strong sense, as a real kind in the world.

Another, more common, way of avoiding appeal to the notion of representation that I have outlined is to substitute for it certain other notions, commonly expressed with the word 'representation', that are also present in scientific explanations. I shall call this family of notions, '*information registration*'.

The core of this family is the technical notion of Shannon *information*. In this sense of 'information', one state X 'gives information about' another state Y, if X is statistically correlated with Y to some relevant degree. Thus if workers' getting off from work in China is correlated with a cock's crowing in Nova Scotia, then each gives information about the other. And it can be said, on the relevant usage, that one type of event 'represents' the other. Clearly, this notion of information, as statistical correlation, is not the same notion as the notion of representation that I am explicating. The latter notion is distinctively associated with a psychology. The former is an all-purpose notion of regular correlation.

Sometimes, this information-theoretic notion is adjoined to further notions to yield an explication of the term 'representation'. For example, for state X to represent state Y, state X must not only be statistically correlated with state Y; instances of state X must also be causally dependent on instances of state Y. Since neither the workers' getting off from work in China nor a cock's crowing in Nova Scotia is causally dependent on the other, this more restrictive notion of representation would not apply.

Perception and many beliefs are causally dependent on what they represent. It is clear, however, that this causal information-theoretic notion is not the same as the notion of representation that I began with. A light bulb's being on can be causally dependent on, and informationally correlated with, the flipping of a switch. But the light's being on does not represent the flipping of the switch. Similarly, the direction of a shadow cast by a rock correlates with and is causally

[5] 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).

dependent on the direction of the sun. But the shadow does not represent the sun in any psychologically distinctive sense. No scientific account of the state of the light bulb takes representing the switch as part of its nature. No scientific account of the direction of the shadow makes any reference to verdicality conditions. Nothing about veridicality conditions and nothing of psychological interest resides in the electric circuitry of the lighting system, or in the interaction between the sun's rays and the rock. One can call the relevant relation 'representation' if one wants. But one is using causal and correlational notions, not a psychologically distinctive notion of representation.

Sometimes, the information-theoretic notion, together with a causal notion, is adjoined to a notion of function, usually biological function. The result is often termed 'representation'. Thus for state X to 'represent' state Y in this sense, X must not only correlate with Y and be causally dependent on Y; X must also have the biological function of correlating with Y. This conjunction of notions would exclude the light bulb's state from representing the switch's being flipped and the rock's shadow from representing the sun's direction. But this conjunction of notions is still vastly more inclusive than the notion of representation that is relevant to characterizing mind. A plant's growing in a certain direction correlates with, is caused by, and functions to correlate with, the direction of the sun. But the plant's states do not represent the direction of the sun in any psychologically distinctive sense. No appeal to veridicality conditions figures systematically in accounts of processes in the plant.

Again, one can call the correlation-causal-functional complex 'representation' if one wants. But one is not using a psychologically distinctive notion of representation. No science explains the plants' growth in terms of its having states with veridicality conditions as aspects of their natures. They are purely biological processes. There is a very broad, highly generic relation between this type of information registration and the psychological notion of representation that is used in perceptual psychology. But the differences are palpable.

Many of the sensitivities of plants and other organisms have been illuminatingly explained in terms of information registration. Many biologists and philosophers use the term 'representation' very broadly to comprise all information registration. Doing so curries interest. It is sexier and gets more headlines and grant money to say that the bacterium represents oxygen and knows how to get away from it, or that the tree sees the sun. But nothing about veridicality

conditions – no representation in a psychologically distinctive sense – enters into the scientific account. Explanations in science that use some notion of information registration do not use the notion in explaining the *formation* and *processing* of states. The notion enters only into background functional explanation – an account of what the biological structures are *for* in the survival of the organism. The basic states and processes of the plant are explained purely in bio-chemical terms or structural or statistical-informational terms that are not in any ordinary sense psychological.

There is, for example, a functional explanation of why a plant's biochemistry takes the form that it does, how the plant's growth is to be understood in evolutionary terms, and what function such growth has – how it was selected for. But no science takes the plant's internal states to have accuracy or truth conditions, with distinctive structures and semantics. Accuracy does not feature in the biology of plant processes. One should not confuse the information-registration sense of 'representation' with the psychological sense of the same word. The extreme breadth of application of the information-registration notions indicates that they differ in significant ways from the notion of representation that has traditionally been thought to be involved in perception, belief, language, and reasoning.

There is nothing in itself wrong with using the term 'representation' in the information-registration way. But such usage adds nothing to explanations in statistical, causal, or functional terms. And it obscures the distinctively psychological kind, *representation*. As I will explain, the more narrowly applicable, distinctively psychological notion of representation is central to a mature scientific enterprise. So substituting information-registration notions for the stricter, psychologically distinctive notion of representation encourages failure to recognize scientifically important distinctions. It also encourages either a romanticism to the effect that the whole living world is psychological or an unearned reductionist view that psychology is 'just' biology.

There is a scientific difference between information registration and psychologically distinctive representation. In the latter, veridicality conditions play a role in *actual scientific explanation*. They play such a role most impressively in perceptual psychology. The point of perceptual psychology is to explain causally how individuals perceive particulars and attributes in the physical environment, and under what conditions individuals fall into perceptual illusions. This form of explanation has been present in psychology since Helmholtz's work in the late 19th century. But it has become the

centerpiece of mathematically rigorous, systematic work that has become a mature science since the early 1970s. The science of *visual* perception has become a more impressive science than many parts of biology, including many parts of neuro-science. Since science is an unsurpassed 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.

Representationally successful perception is perception that is accurate about an environmental subject matter. Perceptual illusion is a mistake about a subject matter. States that are capable of being accurate or inaccurate about a subject matter are postulated in perceptual psychology both as things to be explained and as factors in causal explantions of other perceptual states. Representation, in the sense that I have outlined, lies at the center of a mature rigorous science.

The practice of perceptual psychology contrasts quite markedly with the science of plant biology and the sciences of very simple organisms like bacteria, paramecia, amoebae, and so on. Information registration, in various forms, enters into explanation in these sciences. States with veridicality conditions play no central role. 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.

There are reductionists who believe that the information-registration use of 'representation' can do all the scientific work that the psychologically distinctive notion of representation does. They hold that the only scientifically acceptable notion of representation does apply just as literally to the sensitivity of plants to light and the sensitivity of paramecia to chemical compounds as it does to perceptual states and processes. They may simply regard the 'real' explanation of psychological processing to be entirely at the neuro-biological level. Or they may simply hold that psychology can get by with correlation, causation, and biological function. What appear to be psychological explanations of processing are simply glosses on the biological function of the relevant processes, just as the functional glosses on plant growth are.

Of course, it is an empirical question whether one kind of scientific explanation can be reduced to another. But reductions must be

earned. Explanation in perceptual psychology features veridicality conditions very centrally. It is not acceptable simply to assume such a reduction. Current scientific practice simply does not support such an assumption.

I believe that there are systematic reasons why reduction of representation to any of the types of information registration that figure in science is unlikely. I will not go into detail here. I will just sketch the main idea. The nearest thing to the notions of accuracy and error in the suite of ideas that comprise various types of information registration are notions of fulfillment or failure to fulfill biological function. Success in realizing a biological function is a practical matter – fitness for survival. But truth, accuracy, falsity, and inaccuracy are not practical matters. In principle, biological success could be correlated with inaccuracy, all the way down; biological failure could be correlated with accuracy, all the way down.

Error can contribute to fitness. A rabbit's repeated misperceptions of danger combined with quick-trigger reactions might not only insure against capture. Its expenditure of energy in acting on those misperceptions could make it more adept at fleeing. Conversely, accurate representation can contribute to failure of fitness. Accepting truths about how things really are can lead to disfunctional depression. The practical value of contribution to fitness is simply not the same as the representational value of veridicality.

Of course, contribution to biological fitness and accuracy *do* track one another by and large, especially at the most primitive levels of representation. The close connection between accuracy and fitness, however, does not affect my primary point. The difference between the practical notion of biological function and the notion of veridicality is so fundamental that an explanation that centers on biological (or artifactual) functionality inevitably explains something different from an explanation that centers on representation that involves veridicality conditions. So a purported reduction changes the subject.

In summary, I think it overwhelmingly unlikely that representation can be scientifically reduced to any type of information registration. There is no question that *representation* that sets veridicality conditions and *information registration* are different theoretical notions. I believe that they have irreducibly different explanatory potentials. Explanations of accuracy and of formation of perceptual states capable of accuracy or inaccuracy explain different matters than explanations of contribution to fitness. Perceptual psychology postulates representation as its central kind in its primary scientific

explanations. In philosophizing about a science, the most reasonable starting point is to accept the commitments of the science itself.

I have been making some references to the science of perceptual psychology. Its postulation of representational states indicates that states with veridicality conditions are explanation-grounding kinds or natures. There appears to be no more primitive representational psychological state than perception. Perception is where representational mind begins. I will be trying to explicate commitments of perceptual psychology as a starting point for understanding this beginning.

The first basic point about perception is that it is a psychological state with veridicality conditions as part of its nature.

A second basic point depends on knowing something more about perceptual psychology. I will say a little about the shape of the science. The central problem of perceptual psychology, paradigmatically *visual* psychology, is to explain causally how veridical perceptions and perceptual illusions are formed from sensory input. The science presupposes explanations in physics that connect environmental events with impacts on the sensory receptors. For example, visual psychology assumes the account in optics of how light is propagated from a surface of a certain size, shape, reflectance, and in a certain position, to impacts on the sensory receptors. The optical laws of light, as it projects from a surface to a certain array of frequencies as they strike the retinal receptors, are well understood.

The psychological account has two main forms. One explains a causal chain of states that begins with the registration of the inputs into the visual system – for example, the first effects of the light array on the retinal sensors – and ends with perceptual states that represent the environment as being certain ways. The second form explains how environmental entities reflect light into the retina so as to yield perceptual states. This second form combines the first form with the background assumption of the optics. This second form enables the psychological science to explain accurate perception and illusion. Explanation of this sort can both anticipate general environmental conditions under which perception will be accurate or inaccurate, and explain successful perception and illusion in particular cases.

What makes both forms of explanation difficult and interesting is what is known in the science as the *under-determination problem*.

Visual perception represents particulars and attributes in the environment. But the initial states of the perceptual system are sensory registrations of proximal stimulation. Proximal stimulation is the stimulation closest to the sensory receptors. Such registrations of

light arrays impacting the retina are not perceptions. But these and other registrations are all that the visual system has to go on.[6] Different environmental conditions can produce the same 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.[7] That is, a given registration of proximal stimulation is in itself compatible with many possible perceptual states.

So there are two types of underdetermination that confront the central explanations of perceptual psychology. One is underdetermination of the environmental objects of perception by registration of proximal stimulation. The other is underdetermination of perceptual states by the registration of proximal stimulation.

The *underdetermination problem* is that of answering the following question. How are perceptual states that represent specific particulars and attributes in the environment produced, 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?

The initial registration, or encoding, of proximal stimulation – mainly registration of light – goes through a series of transformations in the visual system that eventuate in perceptual representations that represent entities in three-dimensional space. There is a determinate optical and geometrical solution to the problem of determining how a three dimensional array projects onto a two-dimensional coding of that array. But there is no determinate mathematical solution to how a two-dimensional coding is transformed into a representation of, and as of, a three-dimensional scene. The retinal encodings, together

[6] The visual system utilizes a wider range of input, not just stimuli of the retina. For example, it uses proprioceptive information about the direction and movement of the eyes. The underdetermination problem applies to the wider range as well.

[7] Undetermination 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.

with all other 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, through a series of transformations. These formation laws are distinctively psychological. They systematically cite representational states in terms of the accuracy conditions of the states – for example, a perceptual state that specifies surface or body X as farther away, by such and such distance, from surface or body Y.

The formation laws, and law-like patterns of processing, in effect *privilege* certain possible environmental causes over others. The effect of the privileging is that the under-determining registration of the proximal stimulation triggers, through a series of transformations, a perceptual state that represents exactly one of the many possible environmental causes that are optically compatible with the incoming light and its registration.

The underdetermination of environmental causes by proximal registrations renders the formation of perceptual states subject to error. Illusions occur when abnormal, 'unprivileged' environmental causes produce the same types of registrations of proximal stimulations that are produced by normal, 'privileged' distal causes. These conditions are a topic of the science.[8]

The foregoing discussion of the underdetermination problem grounds my second point about perception. The point is this: formation of perceptual states constitutes a certain type of *objectification*.

Given that a process yields a perceptual state that specifies (has accuracy conditions regarding) environmental entities, formation of perceptual states involves a type of objectification. This objectification is formation of a state that functions to represent a subject matter beyond idiosyncratic features of the individual. The subject matter is the physical environment. The objectification involves a certain removal from the local or idiosyncratic.

Objectification resides in the ways perceptual systems overcome the two forms of underdetermination. The perceptual system distinguishes

[8] Most points made here occur in any mainstream textbook in visual psychology. See Stephen E. Palmer, *Vision Science* (Cambridge, Ma: MIT Press, 2000), 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, 4th edn.). The real science resides in journal articles.

patterns coded in sensory registration that are likely to be adventitious, or idiosyncratic to the perceiver, from patterns that tend to correlate with specific aspects of the environment. When perceptual processing yields representation of the physical environment, it constitutes the relevant *objectification*. The sensory registration is local and idiosyncratic. The perceptual states represent a reality beyond proximal stimulation.

What marks a system or state as *perceptual* is a processing that contrasts registrations of proximal stimulation (and higher-order features of such stimulation), on one hand, with states that specify elements of the environment beyond the sensory receptors. Explanation of the formation of the relevant states benefits from taking the states as representing the environment rather than merely filtering and weighting proximal stimulation so as to make behavioral responses efficient in functioning to respond to environmental pressures.

Objectification in perception is implemented by *perceptual constancies*. Perceptual constancies are perceptual capacities systematically to represent some given particular or attribute as that very particular or attribute under significant variations in registration of proximal stimulation. In a perceptual constancy, a perceptual system can represent some given aspect of the physical environment as that aspect from different *perceptual perspectives*, produced by different proximal stimulation.

For example, *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 perceived particulars as at a given distance and direction, under various types of stimulation deriving from various types of particulars perceived. The sonar systems of bats, whales, and dolphins and the visual systems of numerous animals can localize objects, even if the objects in a given location change many of their properties so as to produce very different proximal stimulation. *Luminance 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. There are numerous other types of perceptual constancies.

Perceptual constancies are marks of objectification. I conjecture that a sensory system is perceptual if and only if the system includes perceptual constancies. This conjecture is not based on conceptual analysis or definition. 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. The idea of the conjecture

is that the central aspect of perceptual systems that makes it necessary to explain formation of their states in terms of representational contents with veridicality conditions is the presence of perspectival capacities inherent in perceptual constancies.

My conjecture is that, in the actual world, in the absence of perceptual constancies, 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. Such explanations do not ascribe perceptual constancies.

I believe that this difference in explanatory strategy 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.

As I have intimated, many species exhibit perceptual constancies. Some arthropods – bees, locusts, and some spiders; most reptiles, amphibians, and fish; and probably all birds and and nearly all 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 lack it. We are fortunate. 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.[9]

Since certain arthropods have perception and since perception involves the most primitive type of representation, these animals have the most primitive type of representational mind. Representational mind begins with bees, spiders, locusts, and preying mantises.

Let me summarize my overall argument. Perception is a natural psychological kind, recognized in rigorous, mature science. Perception is marked by having accuracy conditions as part of the nature of the kind. The formation of perception involves a type of objectification. This objectification is a process that systematically contrasts phenomena that encode proximal stimulation, at various levels of abstraction, and phenomena that represent specific environmental entities. Objectification is marked by processes embedded in exercises of perceptual constancies. These are perceptual capacities to represent some environmental particular or attribute *as* that particular or attribute under a large variety of proximal stimulations, and

[9] See *Origins of Objectivity*, op. cit., 419-421.

from a large variety of corresponding perspectives. The simplest animals that are known to exhibit perceptual constancies are bees, certain spiders, and other arthropods, such as locusts and preying mantises. So they have perceptual capacities. Perceptual capacities constitute the most primitive sort of representational mind. So representational mind begins in the arthropods.

I return to the other mark of mind – consciousness. In our current state of knowledge, I see no point in wrangling over priority between consciousness and representation. I think it better to follow out each of these marks of mind, to see where it leads. We may come to think that our concept of mind straddles importantly different kinds – representation and consciousness. Or we may come to understand some sort of unity between the two primary marks. If someone wants to reserve the term 'mind' purely for systems capable of consciousness, I will not object. I prefer to talk about *representational mind*, or *representational aspects of mind*, and *conscious mind*, or if you insist, *representational psychology*. I think that understanding the main marks of mind or psychology is, at this stage, the main matter.

In any case, it is not a scientific requirement on perception that it be conscious. We know that bees and spiders have perception. We do not know whether they are conscious. Moreover, there is empirical reason to believe that some perception in bees, and in us, is unconscious.[10]

[10] I cite three bodies of empirical evidence, among many others, with a sampling of relevant psychological literature.

(1) 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. Adrian G. Dyer, 'Bumblebees Directly Perceive Variations in the Spectral Quality of Illumination', *Journal of Comparative Physiology* **A 192** (2006), 333–338. There is reason to believe that similar, very primitive color constancies occur at the retinal level in humans–again, well before consciousness is likely to arise. See M.T. Vanleeuwen, C. Joselevitch, and I. Fahrenfort, 'The Contribution of the Outer Retina to Color Constancy', *Visual Neuroscience* **24** (2007), 277–290

(2) There are, in humans and other animals, certain states formed post-retinally, but in the first micro-seconds of visual processing, that seem to involve perceptual constancies, 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. Steven J. Luck,

Moreover, not all consciousness involves perception, or even representation. Awareness of the felt quality of pain (as distinguished from

---

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.

(3) We know that blindsight patients perceive entities, again showing an array of perceptual constancies. It is likely that 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. I discuss these matters in greater depth in *Origins of Objecvtivity*, 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. Umilta 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',

proprioceptive locating of pain) does not require representational content or perceptual constancies.

There may be organisms that feel pain – and hence are conscious – but lack any representational capacities. There may be organisms that have representational capacities, but are never conscious.

We do not know where consciousness begins. We do not know where awareness of pain begins in the animal world. We do not know enough about consciousness even to know how to directly investigate the matter.

But we do know where representational mind begins. It begins with the arthropods.

This result may not make us feel any special kinship with insects. They do not make good pets. They do not love or appreciate us. They lack endearing eyes. But we do have this in common. We are both capable of representing aspects of the physical environment, in a distinctively psychological sense of 'represent'. We both have representational mind. In this respect we are part of a small minority in this universe dominated by rock and fire. And we differ importantly from another minority with which we share life – the plants and other pre-psychological organisms.

TYLER BURGE (burge@ucla.edu) *is Distinguished Professor of Philosophy at University of California, Los Angeles. He has written on philosophy of mind, philosophy of psychology, epistemology, philosophy of language and logic, and the history of philosophy – principally Frege, Descartes, and Kant. He is the author of* Origins of Objectivity *(Oxford, 2010) and three collections of articles:* Truth, Thought, Reason: Essays on Frege *(Oxford, 2005);* Foundations of Mind *(Oxford, 2007); and* Cognition Through Understanding *(Oxford, 2012).*

Neuropsychologia **45** (2007), 3087–3091.