

# Perception, Philosophical Issues about

Austen Clark  
Department of Philosophy  
103 Manchester Hall, U2054  
University of Connecticut  
Storrs, CT 06269-2054 USA

In Lynn Nadel (ed.), *Encyclopedia of Cognitive Science*, London: Macmillan Reference Ltd., 2003. Volume 3, pp. 512-517

**Abstract.** A discussion of three philosophical issues that have arisen within the contemporary scientific study of perceptual processes: the variations and limitations of representational systems, the localization of perceptual function in neural tissue, and the relations to verbal outputs of ordinary citizens introspecting their ordinary minds.

## 1. Central philosophical issues about perception

It might be salutary first to discharge some vapours surrounding the notion of a “philosophical” issue as opposed to all the other sorts of “issue” one might stumble across while exploring the vast domain of cognitive science. Broadly speaking, philosophers work on three kinds of problem. The first consists of all the work one must do to a question before one can say anything sensible about it. The second try to spell out the relations between answers to one set of questions and answers to others. The third, and most advanced, try to reconcile the internal tensions that eventually arise when one tries to spell out the relations between answers to different sets of questions.

The first and third of these kinds of problem are studied almost exclusively within departments of philosophy. Applied to perception, our first kind includes questions such as: Should I believe my senses? Can one prove the existence of the external world? Is it fair to treat the existence of the table in front of me as nothing more or less than a plausible *hypothesis*? And so on. These are classic questions of epistemology, studied within philosophy in a more or less continuous fashion over the past twenty five centuries. They are of immense importance, and well worth studying (see Chisholm 1957). But one must have already made one's peace with them before even starting the enterprise of “cognitive science”—or any science, for that matter—so, given my brief compass, I must set them all aside.

The third kind of question is likewise interior to philosophy. In the attempt to come up with some systematic overview of our intellectual landscape—some synoptic view of how it all fits together—various schools and traditions have arisen over those centuries, and they have their own disagreements with one another. Different people have gone about this project in different ways, and what we have in our third batch are arguments between them: arguments about how to argue, questions about the enterprise of raising questions. Within perception this third category of problem includes such questions as: what sort of logic is required in order to describe the things one seems to see in an illusion? (see Hintikka 1969) Do those things actually exist in some other possible world, or merely seem to? Can they cause things to happen in this world? And so on. These are the sorts of questions that contemporary analytic philosophers spend most of their working hours working on, and while as a philosopher I think such work is the most important work one can possibly do, the enterprise is admittedly an esoteric one, and will appeal to few cognitive scientists.

So this article will confine itself to questions of the second and middle kind: questions that arise over relations between different portions of the scientific study of perception. Within cognitive science itself there are problems about perception that are philosophical problems—problems that arise as soon as one steps back from work on some particular research question, and asks instead: how would the answer to this question cohere, or fail to cohere, with answers to all these other questions that we have already settled, over in this other place? One strives for a pleasing and coherent overview of how it all hangs together; but one finds places where the landscape is not at all settled, peaceable, or pleasing. Instead one finds fault lines, where vast tectonic plates grind remorselessly, deep underground, producing earthquakes, lava flows, gaseous effusions, heat, and noise. These are the philosophical regions. I will identify three: three obvious zones of tectonic conflict within contemporary cognitive approaches to perception.

## 2. Impact of cognitive science on issues about perception

### 2.1. Profits and perils of representation

Representation is the wondrous elixir that makes cognitive science possible: a theoretical notion that is powerful, ubiquitous, intoxicating, and dangerous. One way to appreciate its powers is to arrange the various relations of association between classes of events in a hierarchy of increasing orders of logical complexity. At the bottom we find relations of statistical association, of correlation, and of causation.

Claims of the form “ $x$  is correlated with  $y$ ” are common to all branches of cognitive science and require no special theoretical tools. The relation is fully *extensional*, in that if  $x$  is correlated with  $y$  and  $y = z$ , then indeed  $x$  is correlated with  $z$ . Causal links are one notch up, with some additional content. Another notch gets us to information. Like correlation, talk of information describes a kind of association between ensembles of classes of events, but it is a more complicated kind. It requires a rather robust structure of relations of conditional and a priori probabilities between ensembles of input events and output events. These relations can help one make discriminations that mere causal or correlational talk cannot. They may for example help pick which object is the one perceived among all of those in the causal antecedents of a given perception (see Dretske 1981, 155-68). Causally all those antecedents are of a piece, but the perceptual state carries much more information about some of them than about others. Talk of information can relate classes of events in ways that causal talk cannot.

Trotting out the term “representation” adds yet another order of complexity to our talk. Event  $x$  represents  $y$ , is about  $y$ , says something concerning  $y$ , and is more or less accurate or inaccurate in what it says about  $y$ . What makes this relation of association more complicated is that it proceeds through a semantics: one must provide a semantic interpretation for events at the “representing” end of the relation. That is, what it means to say that  $x$  is a representation is that it has some content “about” some putative object  $y$ , and that such content can be assessed for correctness or incorrectness, accuracy or inaccuracy, truth or falsehood. In order to understand that content one must understand exactly what  $x$  is representing, and in order to do that one must understand the semantics of the system of which  $x$  is a part. Extensionality fails: even though  $x$  represents  $y$  and  $y = z$ ,  $x$  may fail to represent  $z$ . It all depends on the semantics, on how  $x$  represents  $y$ . It might represent its object as  $y$  but not as  $z$ —as water, but not as  $H_2O$ , even though in fact water =  $H_2O$ . The system in question has yet to learn its chemistry.

The notion is intoxicating because it is so powerful. Once we endorse the claim that the objects under investigation themselves employ a system of representation, then suitable tinkering with the details of those systems can explain any behavior one might encounter, or any behavior one pleases. And indeed the cognitive revolution is (arguably) founded on the claim that all mental states are representational; within our theories the elixir can be employed anywhere, at any time, in liberal quantities.

Perception provides a particularly interesting test-case for this platform. Is all perceptual content representational? Is any? What makes perception so interesting in this regard is that its scientific study started at least a century before the cognitive revolution; competitive research traditions, confining themselves to simpler and less powerful relations of association, were already well established when that revolution occurred. Members of the *ancien régime* had accomplished much with the older tools of correlation, causation, and information, and mounted some resistance to the grandiose claims on behalf of the hegemony of representation.

Even today to the questions “Is all perceptual content representational? Is any?” one finds a gamut of opinions ranging from “Yes, all of it” to “No, none of it”. Clearly there are aspects of perceptual experience that do not seem to represent anything: whatever a sensation of green signals, could, it seems, just as well have been signaled by a sensation of blue. So the difference between green and blue is not a difference in what the experience represents. Other qualitative variations likewise lack representational content. The penny on the table presents varying species of ellipse at different times, from different perspectives, even though, throughout one’s travels, one would never judge that penny to be anything other than circular. All the aspects of object constancy—size, shape, color, and so on—provide similar examples. The variations in the character of perceptual experience that are discounted when one achieves object constancy do not represent variations in the object. They are variations in what philosophers call “qualitative” content, and they seem not to represent anything (see Peacocke 1983, 4-26).

The most prominent surviving representatives of the “no, none of it” school are the ecological psychologists and perception-action theorists. The ecological psychologists are wary of the intoxicating charms of the representational elixir, and prefer to cast their theories using relations of association that stop at the level of information pick-up (see Gibson 1979). Perception after all is found in creatures so simple that one hesitates to ascribe them any other cognitive state at all, and perhaps their commerce with the environment can be accurately described without invoking the complexities of semantic interpretation. Such is the promise of perception-action theory.

The thorough-going nay-sayer runs into difficulties, though, when confronted with the problem of perceptual illusion. For at least a millennium the conventional wisdom about illusions is they provide an example of misrepresentation: an error of the senses, a case in which the senses deceive. And error and deception are impossible unless

correctness conditions apply. These states have a content, which in such cases is, sadly, misleading. Other varieties of perception that are less than fully accurate, such as the perception of pictures, of mirror images, or of movies, raise similar issues. All provide cases of “mere appearance” (intentional *inexistence*), in which someone perceives something that seems to be *P*, or seems to perceive something that is *P*, though in fact there is nothing in the vicinity—nothing within the optic array—that is *P*. The representational gambit is particularly hard to resist in such cases. The alternative is either to go metaphysical (to posit some new entity in the situation that is *P* and that is somehow directly perceived) or, heroically, to deny the existence of illusions altogether. Whereas the representational account is straight-forward: the subject is in such episodes representing something to be *P*, but that representation is (for one reason or another) a misrepresentation. What is real is the existence of the representation. What is unreal is the thing represented. We have a real representing of something unreal. This neatly disposes of the ontological qualms caused by mere appearance.

Or, at least it does, if you are happy to accept the hegemony of representation, and allow its governance to extend all the way to the earliest stages of the sensory processes of the simplest creature. This revolutionary program still needs to be empirically substantiated. Doing so is not an easy job. What are the elementary terms, the morphemes and syntax, of these systems of representation? One needs to specify the primitive elements, the rules by which they are combined, how they refer, what predicates they employ, what their truth conditions are. Are they simple, isolated “features” that might be registered in cortical feature maps? Or must even the primitive terms have the relational character of *gestalten* (or of affordances)? Is the primeval form of sensory reference allocentric or egocentric? Are all the predicates innate? Is the comprehension thereof modular? And so on. Plumping for the representation relation adds another degree of freedom and another layer of complexity to one’s theory. While this makes the theory more powerful, the cost of quaffing the representational elixir is that someday the bill comes due: one must pay out all the empirical details of how those representational systems are constituted. There have long been grounds for skepticism on our ability to do this—to have warrant for ascribing *this* system of representation to a creature rather than *that* other, extensionally equivalent one—and the jury is still out. Like any revolution the cognitive revolution makes some bold promises, and sometimes it takes several generations to learn which are lies.

## 2.2 *Revolt of the physiological underlings*

Another region where there are strains within the governing coalition is found in the relations between the top-down, high-level, boss, chief executive psychologists and the bottom-up, low-level, worker, underling, physiologists. This description is deliberately tendentious, but it reflects the temper of early broadsides in the revolution. The boss executives were supposed to run the show, tell us what jobs needed to be done, and give broad functional specifications of the different subsystems of the mechanism they master. The neuro-physiologists, bio-physicists, and neuroscientists gratefully receive these job specifications, get to work, and eventually pass upwards the implementation details of how the lowly neurons cooperate so flawlessly yet mindlessly to produce the wonders of representation. The physiological workers, tinkering in the labs, would be lost without such guidance from their mindful masters on high. The “autonomy” of levels of explanation guaranteed that the offices of the executives would never be invaded by the grubby workers. It even seemed possible to make room in the executive suite for some philosophers, and some philosophers gratefully signed on, and moved in.

This pacific vision of corporate harmony has been discomfited by the pressure of empirical discoveries. What are we to make of “feature detectors” in the primary visual cortex, or of the subsequent “feature maps” in secondary areas? (Treisman 1993) How are we to understand the distinction between what *v.* where channels in visual processing, or the importance of synchronized 40 hertz oscillations? (Crick and Koch 1990) These rumblings and pressures, gurglings and belches, definitely proceed from the bottom up, and they rudely interrupt the a priori musings in the executive suite. Suddenly the physiological underlings seem to be providing job specifications for the psychologists, instead of the other way round. Such developments indicate the need to rethink relations between sensory contents and physiological implementation (see Clark 2000).

It would be good to have a clear view of those relations: of relations between perceptual states and neural states. For example, might we someday find a “colour perception center” where perceptions of colour are localized? The question poses a dilemma, in that both alternatives seem equally bad. It remains incredible to imagine that we someday point to some region of Joe’s neuroanatomy and say, truthfully, “process *x*, going on in there, is identical to Joe’s sensation of red”. How do goings-on in that region get promoted to the status of “sensation of Joe”? And how could we explain the connection

between goings-on in that region and the fact that the sensation is a sensation of red, and not of green? This is one version of the “explanatory gap” (Levine 1983) and it has yet to be plugged. But alternatives to localization are equally hard to fathom. How could the experience of seeing a red patch be distributed across chunks of neural tissue? How are the chunks coordinated? What glue melds them into one experience, or into an experience of one patch, and not two? (see Crick and Koch 1990)

Other questions under this heading are equally in need of recalibration. The connected claims for the “autonomy” of “levels of explanation”, and for psychology as a “special science” (Fodor 1979, 1-26), seem increasingly irrelevant to the scientific study of perception. Can one consistently hold such views while maintaining that all properties are physical properties? (Kim 1993). Some other keywords to keep handy for future developments: psycho-physical laws (Davidson 1970), physiological reductionism (Bickle 1998), and mind-brain identity (Hill 1991; Warner & Szubka 1994).

### 2.3 *Triumphs and disasters of public relations*

A final locus of seismic conflict for contemporary cognitive approaches to perception lies on the tense border between it and the legions of ordinary language and common sense. The scientific study of perception is blessed or cursed with the fact that many of the processes so arduously described in its theories seem also to be open to direct observation and immediate access by the ordinary Joe. After all, Joe sees, touches, tastes, smells, and hears; is often aware of doing such; can say when he does such; and can describe something of what it’s like when he does. What are those processes but the very ones described by theories of perception? So the theories seem to enlighten or encroach upon a territory already staked out and claimed within the provenance of common sense. Practitioners are therefore forced to stipulate relations between what they say and what the legions of common sense say. This job is quite literally a job of public relations, and like any public relations job it is delicate, essential, and potentially explosive. A single misstep or indiscretion could prove fatal to the entire enterprise.

It is under this heading that one must place all the philosophical arguments relying on ordinary intuitions about what it’s like to see red, about the “qualitative character” or “qualia” of seeing red, and about the possibility that what it’s like for me to see red is qualitatively identical to what it’s like for you to see green (the so-called “inverted spectra” arguments; see Block 1980). All such arguments rely on

truth claims based on common sense intuition. This does not imply that the arguments can be entirely dismissed—claims based on common sense might occasionally be true—but simply that the task of figuring out a response to them is similar to the task of negotiating a relationship with a public whose folkways and intuitions differ from one’s own.

Answers here run the gamut from ignoring ordinary Joe altogether (either denying that perceptual experience has any contact at all with the states and processes hypothesized by these theories, or admitting that there is some contact, but despairing of any useful information arising from introspective methods) to embracing his verbal outputs as constituting the canon for “heterophenomenology” (Dennett 1991). At one end of the spectrum we have the public relations triumph of scientific vindication of all the intuitions of dear old Auntie (Fodor 1985)—the customer is always right!—while at the other end we contemplate the consequences of the corporate PR man announcing, gleefully, that Auntie and all the other putative customers have no thoughts, no desires, no beliefs, and no existence (Churchland 1979).

The problem can sometimes move directly into the lab, and does so when subjects for experiments on the cognitive unconscious walk in the door. What are we to make of the appearances (or lack of appearances!) of first-person introspective access to the states and processes hypothesized by theories of perception? For example, blindsight is problematic largely because one relies on the truthfulness of DB’s testimony that he does not see the *X*, is not aware of the *X*, even though his pointing (and other behavioral measures) would normally be taken to show that he sees it perfectly well (Weiskrantz 1997). But why should we take DB at his word when he says, sincerely, that he does *not* see? And why should one adopt the common sense assumption that seeing something implies consciousness of it, or, alternatively, consciousness of seeing it? Resolution of these issues requires nothing more or less than settling the old problem of the scope and limits of introspective methods.

Some distinctions that have gradually become clear in the philosophical literature may be of help. For example, qualia—the very paradigm of entities whose essence is fully revealed under the glare of introspective consciousness—turn out to be ambiguous. In one sense qualia are the qualities attributed to appearances—the qualities of colour, taste, smell, etc. that things seem to have—and in another sense qualia are the properties of sensation: the properties of sensory states in virtue of which things appear as they do (Clark 2000, 1-11). Neither notion should be confused with “what it is like” to have a

given sensory state. The latter phrase, made popular by Nagel (1979), was meant to pick out all and only the *conscious* mental states, and if one admits the possibility of sensory states of which one is not conscious, then the qualia of a sensory state may sometimes diverge from “what it is like” to have that state (Lycan 1996). The term “conscious” itself has manifold ambiguities; in one sense it applies to any creature that is awake and sentient; in another sense it applies only to mental states of which one is introspectively aware (Rosenthal 1997). In any case it is very clear that consciousness and sentience are not the same thing. One can study sensory processes and the qualitative character of sensory states without necessarily committing oneself to some verdict about what it is like to have those states.

Such Balkanization is likely to continue, and its logical endpoint would be a continuum of states, starting at some that are clearly unconscious and insentient, and ending with some that are, in the fullest sense, conscious sensory states. In between we find a vast series of intermediaries, of varying orders of organization, encompassing all the transition zones between unconscious, semi-conscious, and partially conscious, and capturing all the pathological and paradoxical breakdowns that might occur along the way. Once that order is empirically described, and relations between all the neighboring points within it are made clear, it becomes a matter of indifference where one draws the line between cases in which some ordinary language term applies and cases in which it does not. All the facts upon which such a verdict rests would be laid upon the table, and the only questions that remain would be verbal.

### 3. Relevance of philosophy of perception to cognitive science

The three philosophical issues about perception outlined above are rarely found in pure forms, as isolated and distinct dilemmas about representation, about neural implementation, or about common sense introspective access. More typically, and more potently, in the outstanding controversies of the day all three themes are typically found intermingled, in a maximally confounding combination. So for example, the puzzles over temporal anomalies of perception and the phi phenomena (Dennett and Kinsbourne 1992) require all three ingredients to reach critical mass. We need assumptions about how time is represented, how those representations are instantiated in the nervous system, and how ordinary subjects access those representations when they see a moving light. “Filling in” likewise gains the critical mass of a dilemma only when one imports assumptions from all three of our philosophical regions: assumptions

about spatial representations, about their neural implementation, and about introspective access thereof (Pessoa, Thompson, and Noë 1998). Blindsight (Weiskrantz 1997) and the various “binding problems” (Crick and Koch 1990; Treisman 1993) may seem to be almost pure neurophysiological issues, but on closer examination what makes them controversial, what generates all the heat, is some combination of assumptions crossing our philosophical fault-lines. Lots of energy builds up at such places, and when they let go, the results can be exciting.

### References

- Bickle, John (1998). *Psychoneural Reduction: The New Wave*. Cambridge, Massachusetts: MIT Press.
- Block, Ned (1980). Troubles with functionalism. In Ned Block (ed.), *Readings in the Philosophy of Psychology*, i. Cambridge, Massachusetts: Harvard University Press, 268-305.
- Chisholm, Roderick (1957). *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- Churchland, Paul M. (1979). *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press.
- Clark, Austen (2000). *A Theory of Sentience*. Oxford: Oxford University Press.
- Crick, Francis and Koch, Christof (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences* 2: 263-75.
- Davidson, Donald (1970). Mental events. In Lawrence Foster and J. W. Swanson (Eds.) *Experience and Theory*. Amherst: University of Massachusetts Press, 79-101.
- Dennett, D. C. (1991). *Consciousness Explained*. Boston: Little Brown.
- Dennett, D. C. and Kinsbourne, M. (1992). Time and the observer: the where and when of consciousness in the brain. *Behavioral and Brain Sciences*, 15: 183-247.
- Dretske, Fred (1981). *Knowledge and the Flow of Information*. Cambridge, Massachusetts: MIT Press.
- Fodor, Jerry A. (1979) *The Language of Thought*. Cambridge, Mass.: Harvard University Press.
- Fodor, Jerry A. (1985). Fodor's guide to mental representation. *Mind* 94: 76-100.
- Gibson, James J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Hill, Christopher S. (1991). *Sensations: A Defense of Type Materialism*. Cambridge: Cambridge University Press.
- Hintikka, K. J. J. (1969). On the logic of perception. In his *Models for Modalities: Selected Essays*. Dordrecht: Dordrecht Reidel, 151-83.
- Kim, Jaegwon (1993). *Supervenience and Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press.
- Levine, Joseph (1983). Materialism and qualia: the explanatory gap. *Pacific Philosophical Quarterly* 64: 354-61.
- Lycan, William C. (1996). *Consciousness and Experience*. Cambridge, Massachusetts: MIT Press.
- Nagel, Thomas (1979). What is it like to be a bat? In his *Mortal Questions*. Cambridge: Cambridge University Press, 165-180.
- Peacocke, Christopher (1983). *Sense and Content: Experience, Thought, and their Relations*. Oxford: Clarendon Press.
- Pessoa, Luiz, Thompson, Evan, and Noë, Alva (1998). Finding out about filling in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences* 21: 723-802.
- Rosenthal, David (1997). A theory of consciousness. In Ned Block, Owen Flanagan, and Güven Güzeldere (eds). *The Nature of Consciousness: Philosophical Debates*. Cambridge, Massachusetts: MIT Press, 729-53.
- Treisman, Anne (1993). The perception of features and objects. In *Attention: Selection, Awareness, and Control: A Tribute to Donald Broadbent*. Edited by A. Baddeley and L. Weiskrantz. Oxford: Clarendon Press, 5-35.
- Warner, Richard and Szubka, Tadeusz (eds.) (1994). *The Mind-Body Problem: A Guide to the Current Debate*. Oxford: Blackwell.
- Weiskrantz, Lawrence (1997). *Consciousness Lost and Found*. Oxford: Oxford University Press.

## Glossary

**binding problem.** The problem of explaining how the elements of a distributed representation--a representation whose elements are spread out over space and time--can serve together as one representation, or can serve to represent one thing.

**explanatory gap.** The difficulty a physicalist has in explaining why a particular neural state has the qualitative character that it has; for example, why it is a sensation of red, and not of (say) green. The problem lies in ruling out the possibility that that very same neural state could just as well have a different qualitative character. See also spectrum inversion.

**extensionality.** A property of systems of representation in which names (or other terms) that refer to the same thing can be substituted for one another without changing the truth values of representations in which those terms occur.

**feature.** A distinct dimension of discriminability of stimuli. For example, among visual stimuli these would include local contrast, brightness, hue, saturation, local contour, orientation, size, relative motion, texture, etc. Alternatively, a particular value on such a dimension: a particular degree of contrast, brightness, hue, saturation, etc.

**feature map.** A portion of cortex, organized topographically, in which activation of cells of some particular physiological kind *K* seem to register variations in one particular sensory feature.

**instantiation.** The relation between a property and its particular instances. In particular: the relation between a functional specification, a specification of a job to be done, and the particular mechanism that carries out that job. Thought to be a useful way of describing the relation between psychological properties and neural states.

**object constancy.** Constancy in the properties an object is judged to have. More narrowly: constancy in judgements about some particular property of an object, such as its size, shape, colour, etc.

**qualia, qualitative character.** Generically, dimensions of variability in the appearances of things. More narrowly: dimensions of variability in appearance that are not correlated with changes in any objective properties. Even more narrowly: the properties of

sensory states in virtue of which the appearances of things vary in this way.

**spectrum inversion.** The possibility that one person's sensation of red has the qualitative character of another person's sensation of green, and that the switch is sufficiently systematic that this difference is undetectable. Used to argue that qualitative character cannot be defined functionally.