PERCEPTION

Cambridge Handbook to Cognitive Science

Casey O'Callaghan

July 24, 2007

word count: 7956

Suppose you see a ripe tomato drop from above and strike the pavement with a resounding splat. As it captures your attention, you experience the sights and sounds of the tomato and its collision with the ground. You might simply observe the tomato to fall. But you might take a step out of its way or reach out to interrupt its fall. Later, you can imagine that tomato, the sound of its demise, and the mess -- how it was all there before you to experience, attend to, and act upon.

Perception drives discussion in philosophy and the cognitive sciences because it forms our most intimate sort of acquaintance with the world. In perception, the world appears before us as available to our thoughts and susceptible to our deeds. What we perceive shapes our thinking and guides our action. The variety and the flux are impressive.

Philosophical work on perception traditionally concerns whether perceptual acquaintance with things in the world is compatible with the possibility of illusions and hallucinations. Given that you might have been hallucinating the tomato incident, how are you ever acquainted with real tomatoes?

The contemporary cognitive science of perception attempts to understand perceiving in naturalistic terms. Cognitive science aims to explain the *processes* and *mechanisms* by which perceiving takes place in organisms understood as biological systems. The objective is to describe and explain how a creature accomplishes the feat of perceiving given constraints imposed by its physiology, environment, and goals. How do your body and brain, which are made of cells and neural tissue, ground your awareness of the sights and sounds of a crashing tomato?

In this essay, I first introduce the traditional philosophical problem of perception, which concerns whether our sense of perceptual awareness survives arguments from illusion and hallucination. Though it may seem distant from empirical concerns, this problem holds important lessons about the nature of perception and provides the conceptual backdrop to contemporary discussions about what it is to perceive. I next turn to empirically motivated theoretical issues about perception. These concern alternative ways to understand the processes and mechanisms involved in perceiving, the relationships of perception to other forms of cognition and to action, and the variety among things we are able to perceive. In light of recent concerns, I conclude with a discussion of the role of phenomenology in perceptual theorizing.

The philosophical problem of perception

Philosophical worries over perception traditionally stem from a catalog of real and imagined illusions and hallucinations. Sometimes things are not as they appear. For instance, a straight stick half submerged looks bent. The Necker cube illusion involves a flat figure that looks three-dimensional. In the McGurk effect, the phoneme /ba/ sounds like /da/ when viewing a speaker pronounce /ga/. Such illusions involve perceiving a thing while misperceiving some of its features. Hallucinations, however, involve complete figments of experience, as when Hamlet seems to see a dagger, tinnitus sufferers hear ringing, or a lifelike experience turns out to be a dream.

Arguments from illusion and hallucination target the intuitive or naïve view underwritten by perception's phenomenology and aim to prove that perceptual awareness is not a simple, direct relationship to mind-independent entities. If, in some perceptual experience, the world does not match how it seems to the subject, then that experience does not just consist in unmediated openness to things and features in the environment. How, then, since nothing subjective distinguishes illusory from veridical experience, could even accurate experience involve direct awareness of the world? If the objects of illusory and hallucinatory perceptual experiences are not ordinary things and qualities (like tomatoes, colors, and sounds), and if nothing subjective suggests that you are aware of something different in illusion or hallucination than in genuine perception, then perception, too, might be something other than intimate, unmediated acquaintance with extra-mental items.

Philosophical theories of perception attempt to reconcile the phenomenology of perceptual experience with misperception. They are constrained on one hand by the relational or world-involving character of perceptual experience and on the other by the possibility of undetectable illusion and hallucination. It is useful to classify them according to whether the objects of perception, if any, are internal or external (mind-dependent or mind-independent), and according to whether perceptual awareness of things in the world, if any, is mediated by awareness of something else or not (whether it is indirect or direct). The possibility of misperception forces us to take a stand on what sorts of things we are aware of when we perceive, and on that in virtue of which we are aware of them.

To preserve the intuition that you experience the same kind of thing in illusory and veridical perceptual experience, and to preserve experiential contact or acquaintance with the objects of perception, *sense-datum* theories deny that the immediate objects of perceptual experience are ordinary tables and chairs. Instead, they are experience-dependent items of sense that are present to one equally in veridical and delusive experience (see, e.g., Robinson 1994; Foster 2000). But citing private sense data does not suffice to capture how the experience seems. For that, we must mention ordinary public objects and qualities (Strawson 1979). According to sense-datum theories, if one ever perceives anything external and public, one perceives it by virtue of experiencing internal or private sense data (see Jackson 1977). Perception thus is mediated by awareness of internal features of sensation.

Disjunctive theories respond that perception requires contact with ordinary, mind-independent objects, but deny that genuine perceptual experiences and subjectively indistinguishable hallucinations belong to a common psychological kind (see Hinton 1973; McDowell 1982, 1994; Martin 2002, 2004). Delusive states differ from perceptual experiences in lacking a constitutive, subject-independent object. Perception thus is direct because it requires no awareness of internal mediators. According to disjunctive theories, however, subjects are at best in a position to characterize their experiences as, for example, *indistinguishable from when I see a tomato smash*.

Intentional theories, on the other hand, want it both ways. They attempt to capture both the impression that one enjoys perceptual awareness of the external world and the intuitive possibility that subjectively indistinguishable illusory and veridical experiences belong to some common explanatory psychological kind. According to intentionalism, perceiving, like believing and desiring, involves representing things to be a certain way (see Harman 1990; Tye 1995, 2000; Lycan 1996; Byrne 2001). Perceptual states thus have *content* corresponding to what they are about, in the sense in which newspapers, in contrast to buckets, have contents (Siegel 2005). Such states thus may or may not be satisfied depending upon how the world is -- perceiving, like believing, is a psychological attitude with accuracy conditions. It is, however, customary to reserve 'perceive', as we reserve 'know' and 'regret', but not 'believe', for success. Perception plausibly is *factive* in that one cannot genuinely perceive that which is not the case, just as one cannot know or regret that which is false. Like belief, however, *perceptual experience* can go wrong. It might seem to you as if you see a cow even if there is no cow. Your perceptual experience might seem *as of* a cow, though you perceive a cow only if the cow unaccidentally causes your experience. Intentionalists embrace the possibility of misperception while capturing how perceiving seems to involve awareness of things independent from oneself. Illusory and veridical experiences share representational *content*, but among the *objects* of perception are ordinary external things and features. Awareness of public objects is mediated in the sense that it requires representing those objects, but this need not be mediation by some entity of which one is aware. Perceptual representation need not be like viewing a picture. It is more like being the picture.

Since perceptual experiences, as well as beliefs and judgments, can be characterized in terms of their content, what distinguishes perceptual states from other contentful cognitive states? First, perceptual experiences have vivid phenomenology in which things appear as present before you. Additionally, perceptual experiences count as sources and reasons, which bear distinctive evidential weight, for judgments and beliefs. Experiences, however, need not compel belief or judgment. One might withhold commitment to the truth of what one experiences, and need not even acquire a disposition to judge or so commit (Johnston 2006). Conversely, knowing that you are experiencing an illusion may do nothing to change your perceptual experience. Finally, judging and having beliefs are straightforwardly conceptual achievements that require possessing concepts corresponding to what one believes. But, intuitively, people and animals perceive things for which they do not possess concepts: seeing a horseshoe does not require possessing the concept of a horseshoe. Perceiving, then, is not a form of judging or believing. It is an attitude marked by its phenomenology and functional role.

Their contents do, however, illuminate how perceptual experiences differ from mere sensations and from any purely qualitative features intrinsic to experiences. Conscious sensations, such as those associated with pains, nausea, and dizziness, feel private. Such experiences do not seem to involve awareness of something independent from oneself, and sensations themselves do not have accuracy conditions. The possibility of misperceiving a pain, or of nausea existing unperceived, is not compelling. Illusory dizziness is just plain dizziness, and unfelt dizziness is no dizziness at all. Many philosophers hold that sense perception involves purely qualitative sensations, or *qualia*, intrinsic to one's experience (Nagel 1974; Peacocke 1983; Block 1990; see Tye 2003). Though it is a matter of controversy whether any such features exist, they are at least quite difficult to discern (Dennett 1987, 1991; Harman 1990; Dretske 1995; Tye 1995, 2000; Loar 2003). Due to what has been dubbed the *diaphanousness* or *transparency* of perceptual experience, attempts to discern features intrinsic to one's perceptual experience nearly always deliver features, such as colors and shapes, of things one seems to perceive. Any qualia we sense strike us as qualities of things we perceive (cf. Block 1996; Chalmers 2004, 2006).

The mechanisms of perception

Perceptual experience may seem effortless, automatic, and directly responsive to the world. It nevertheless requires a complex battery of subperceptual processes. Sensory stimulation occurs when the environment disturbs a sensory surface, such as the retina, tympanum, skin, olfactory epithelium, or tongue. Receptive surfaces *transduce* chemical, mechanical, or electromagnetic energy into neuroelectrical signals that initiate further sensory and subperceptual processes. How does the disturbance of a sensory surface come to regulate action, impact thought, and stimulate vivid experiences in which you seem to be aware of a richly detailed world independent from yourself? One facet of this puzzle is how processes initiated at the interface between the environment and your sensory surfaces recover aspects of the world just on the basis of retinal, tympanic, or dermal activity. Recent cognitive science attempts to illuminate how the seemingly magical feat is accomplished.

Why is this such a puzzle? In vision, the image projected by the lens of the eye upon the retina is quite different from what we see. The image is two-dimensional, and it is inverted. Due to constant eye movements (saccades and microsaccades that occur up to 60 times per second) the image moves continuously relative to the retina. Rod and cone receptors are distributed unevenly, while image information is lost at the "blind spot" where the optic nerve departs the retina, though no evidence of an experiential gap exists. And a different image strikes each of the two retinas. In audition, air pressure fluctuations set off intricate vibration patterns at the two eardrums. This leads ultimately to a spatial auditory experience comprising discrete sound streams characterized by discernible audible attributes. In olfaction, complex mixtures of chemical compounds cause a huge array of experiences of recognizable smells. The difficult empirical question is: What are the mechanisms by which stimulation of sensory surfaces and sensory transduction lead to full-blooded perceptual awareness?

One might think that after such questions have been answered, we have explained all there is to understand about perceiving. Cognitive scientists and philosophers, however, recently have engaged in heated debate over exactly what these processes tell us about what it is to perceive. On one hand, we might view the evolving states and processes that occur "inside the head" as tantamount to *constituting* a perceiving subject. On the other hand, we might think that although it reveals a critical part of what *enables* perception, citing internal processes as such misses some essential aspect of what it is to be a perceiver. As I hope to make clear in the following three sections, how researchers view the role of activity subsequent to the event of sensory transduction increasingly marks an important theoretical rift concerning the nature of perception. Some explanation is in order.

Perceiving as information processing

The mainstream of cognitive science understands perception as an *information processing* problem. How does one construct a representation of a complex environment, which meets the demands of thought and action, from impoverished stimulation? Viewed as such, developing a theory of perceiving requires discovering the mechanisms responsible for our capacity to extract useful information from sensory stimulation and to present it in a form that suits the needs of other cognitive functions.

Early theorists underestimated how demanding the task is in two respects. First, they construed perceiving as discerning or grasping the significance of a sensory image of which one is already aware. But even consciously sensing an image requires extracting features such as color, shape, and motion from the light. Illumination upon the retina, for instance, is a product of ambient light and surface reflectance, and thus underdetermines color experience. The image theory begs an essential question about vision by assuming awareness of such features without explaining it. What is required is a story about how information is extracted from light to recover even the most basic sensible features.

Second, they took perception, in contrast to sensation, to be continuous with higher cognition and intelligence. Perceiving, understood as discerning or grasping an image's significance, deploys general-purpose cognitive capacities that are involved in many varieties of thinking and reasoning but are not peculiar to perceiving. Lewis (1966: 357) describes one version of the picture like this: "Those in the traditions of British empiricism and introspectionist psychology hold that the content of visual experience is a sensuously given mosaic of color spots, together with a mass of interpretive judgments injected by the subject." General-purpose cognitive capabilities, however, do not suffice to explain how information concerning public objects and things in the world is extracted exclusively from a private image. One might suggest that perceptions concerning, for instance, size and distance result from mechanisms such as learning, association, inference, or intuition. But it is quite unclear how general-purpose cognitive capacities such as associative learning or inference by themselves could yield a robust perceptual judgment concerning three-dimensional arrangements of surfaces and objects if one only ever has immediate acquaintance with a two-dimensional color spot mosaic.

Suppose you take vision to involve consciously sensing a two-dimensional image projected upon the retina, as did Descartes (1637/1988: 64). This two-dimensional image by itself drastically underdetermines the scene. First, the projection of a scene is geometrically consistent with an infinite number of three-dimensional arrangements. A given region may correspond to something small, nearby, square, and at an angle, or it may correspond to something large, distant, trapezoidal, and oblique. A colored image itself determines neither size, distance, shape, nor orientation. Second, the image, marked by color discontinuities, is not clearly carved into units such as surfaces or objects. It thus fails to account for a scene's discernible segmentation into bounded regions with three-dimensional contours and shadows. This deficit is somewhat difficult to appreciate given our facility with pictorial representations, but it is pressing. The capacity to resolve such dramatic ambiguity must be accommodated by any theory of the mechanisms of perceiving.

In a step toward resolving these puzzles, and toward the contemporary approach to perception, Helmholtz (1867/1910) suggested that vision involves a series of *unconscious inferences*. Since such inferences are based upon sensations that possess phenomenal features, and since it relies on a general, associative model of inference, Helmholtz's view exhibits the failings characteristic of early theories (see Hatfield 2002). However, the view according to which perceiving involves inference-like transitions inaccessible to conscious experience anticipates what Fodor and Pylyshyn (1981) call the "establishment view" of perceiving, or what I will call the "received view". It prefigures contemporary accounts according to which *subpersonal* sensory processes fuel unconscious transitions akin to a form of deductive or inductive inference.

Contemporary theories differ in two critical respects from early image and sensation-based theories. First, they recognize that even formulating a color mosaic requires extracting information from light and thus do not take for granted that consciously accessible images or sensations ground perceptual inferences. Second, specialized subsystems, and not a general rational or cognitive faculty, conduct perceptual inferential processing. To transform light information into full-fledged perceptual representations of public objects and features requires representing specific kinds of information at each stage, and requires specialized rules and assumptions to guide the transitions. Perceptual systems are, to a significant degree, automatic and unaffected by one's beliefs and reasoning. Though they furnish materials for thought and action, they are to a great extent *modular* (Fodor 1983; Pylyshyn 1984, 1999). Since deploying general-purpose problem-solving strategies would be radically inefficient given the complexity of the task, perceptual capacities thus forfeit the generality of those proposed by Helmholtz and other early theorists. Perceiving, on contemporary views, may involve processes resembling judgment and reasoning, but given the degree of specialization in perceptual systems, perception is more automatic than classic views supposed. Whether 'inference' is understood literally (Rock 1983) or less-than-literally is partly terminological. Clearly, perceptual processes are not *ordinary* inferences. They are neither conscious nor deliberate. They are suited to a particular kind of task, do not generalize to all varieties of information, and may incorporate little outside information. Hatfield (2002) suggests, nonetheless, that explaining perception in terms of unconscious inferences requires justifying why the cognitive machinery should be understood as conducting inferences, or reaching conclusions guided by premises or evidence. This requires establishing that subpersonal states of perceptual systems have *content* -- that they represent, for example, intensity values, visual angle, and distance -- and implement principles or rules for deriving one from the other. We might also wonder why *conclusions* of such inferences lead or correspond to perceptual *experiences*. Even so, it is tempting to view what takes place in perception as a regimented progression of transitions from early sensory states that bear information about stimulation to later states that represent features of one's environment. Indeed, the received view is that perception implicates information processing characteristic of representational systems (Fodor 1975; Pylyshyn 1984). Perceiving, on this account, is constructing a representation of one's environment from impoverished sensory information.

Gibson and Marr: direct pickup and computation

The view that perception is mediated by unconscious inferences from sensory stimulation is not uncontroversial. Gibson (1979) famously suggested that visual perception involves the *direct pickup* of information about one's environment that is present in the ambient light. Perceiving, according to this account, is not the achievement of representing, augmenting, and transforming the impoverished information of the senses. There is no need, Gibson suggests, to infer or otherwise intelligently construct a rich internal description or representation of the world. The key, instead, is that all the information concerning features that matter to the creature is present in the light that reaches the eye from all directions, the *ambient optical array*. Resolving ambiguities in sensory information requires appreciating that such information is *dynamic* -- sensory stimulation changes over time as a creature negotiates its shifting environment. Since the light that reaches the eye is determined by illumination, the surfaces that reflect and generate light, and a creature's position in the environment, the structure of the ambient optical array changes as a function of illumination and of the movements of both the objects and the animal. The resulting patterns of change, or *optic flow*, contain information about objects and features in the environment. For instance, illumination changes cause relatively uniform changes across the optic array. When an object moves, the pattern of optic flow differs from when the creature moves. The former produces a local, relative change to an otherwise static optic array. The latter produces a distinctive global pattern of optic flow. Walking forward, for instance, creates flow outward from a vanishing point. Perceiving thus involves detecting, in the changing optic array over time as one negotiates the environment, information about invariant properties that correspond lawfully to objects and features. The world, available directly to be perceived, mitigates the need for representations. The world serves as its own model to be explored. According to Gibson, subpersonal processes enable perceiving. But, first, such subpersonal processes should not be understood as intelligently constructing representations from impoverished, static sensory information; second, perceiving is not a subpersonal task. Perceiving is an achievement of nothing less than the creature.

Marr's seminal work *Vision* (1982) begins to reconcile these divergent ideas. It nonetheless belongs squarely with the establishment. Marr's approach likens perception to deriving representations of the world from sensory images according to rule-like constraints. According to Marr, Gibson's insight was to recognize that the senses are "channels for perception" of the outside world, and not simply sources for sensations that fuel cognitive processes. Perception of the environment must be understood in terms of the detection of invariant properties of sensory stimulation through movement and time (1982: 29). Marr, however, claims that, framed as such, "direct pickup" *is* an information processing task that must be carried out by our perceptual systems and that Gibson drastically underestimates its difficulty (1982: 30). Marr thereby acknowledges Gibson's insight that the ambient optic array provides important information concerning invariants, and thus about the arrangement of the visible environment, but Marr proposes an account of how this information is extracted from retinal stimulation and represented in a useful form. Marr thus understands perception in traditional terms: "Vision is a *process* that produces from images of the external world a description that is useful to the viewer and not cluttered by irrelevant information" (Marr and Nishihara 1978: 269).

Marr's innovation is the framework he proposes for understanding perception in *computational* terms. He proposes that explaining the information processing that takes place in perception requires understanding it at each of three levels of abstraction. First and foremost, the *task* or *computational problem* of perceiving is defined by a mapping from information about sensory stimulation to information about the environment that answers to a creature's needs. Second, the level of the *algorithm* (or software) addresses the format in which to represent both sorts of information and articulates a specific solution or detailed strategy concerning how to transform one into another. Finally, the level of *implementation* (or hardware) concerns how physiological processes in the brain realize the computational algorithm.

What distinguishes Marr's account from previous theories is that, to perform the task, so defined, the visual system employs processing strategies that exhibit a grasp upon the *natural physical constraints* governing the sources of sensory stimulation. Since stimulation underdetermines its source, perceiving invokes subpersonal rules that are intelligible only as embodying general principles regarding one's environment. Assumptions about the natural world, concerning, for example, how scene geometry projects to the retinal image or that rigid bodies correspond to patterns of stimulation, govern how visual processes transform representations of retinal stimulation into full-fledged representations of the visual world. Patterns of luminance intensity values are converted into representations of one's visual environment through the help of built-in rules concerning how illumination, scene geometry, surface reflectance, and viewpoint determine luminance (the ambient optic array). A sharp luminance gradient, for instance, corresponds to an edge and thus forms the basis of an edge representation. Perceptual systems resolve the radical ambiguity in sensory stimulation through processing strategies that exploit assumptions concerning its relation to the natural world.

The computational process Marr describes proceeds in stages. From retinal luminance values, a *primal sketch* representing *edges* and *blobs* is computed by detecting sharp intensity discontinuities. From the primal sketch, a *2 1/2-D sketch* encodes information about *visible surfaces*. It represents the contours and arrangement of surfaces in a viewer-centered or egocentric framework that includes information about depth, orientation, and surface discontinuities (1982: 277). The 2 1/2-D sketch depends upon a number of image characteristics and natural constraints. For instance, assuming stereoscopic disparities stem from a common physical source yields distance information; that illumination patterns are generated by rigid bodies yields physical structure from motion and optical flow; that surface elements are uniform in size and distribution yields surface texture from luminance patterns (1982: 267). Finally, a *3-D model*, a detailed description of a scene's three-dimensional shapes, meets the needs for object recognition. The 3-D model comprises primitive volumetric shapes (such as cylinders, spheres, cubes, or cones) assembled with increasing detail to recover the specific geometric structure of the objects in a scene. Object identification might then invoke higher-level cognitive processes, such as pattern recognition and memory, which are beyond the scope of perception.

The information-processing paradigm understands perceiving as transforming sensory information into increasingly rich representations of one's environment. Steps in this process, understood as computations, take place according to algorithms that amount to strategies for interpreting the environmental significance of sensory stimulation. If perceiving culminates with the experience of a visual scene populated by volumes, colors, and shadows, then perceptual systems, guided by natural constraints, must extract such information from retinal clues and build it into a consciously accessible representation. Marr's computational approach to vision, which distinguishes the overall task of vision from the algorithms for its solution and from its neurophysiological implementation, exemplifies the predominant contemporary approach to perception in cognitive science.

Representing and enacting

According to the received approach, perceiving is tantamount to representing. It principally involves reconstructing a representation of the immediate environment from the noisy, ambiguous stimulation of sensory surfaces. Thus all perceptual awareness is, in one sense, mediated by representations derived at the subpersonal level from sensory stimulation.

Understanding perception in such terms has come under fire from a growing antiestablishment.

In the first place, vision may not require constructing a detailed representation of a scene, since we may see far less at any given moment than we take ourselves to see. For instance, it may seem that you are seeing all of the words on this page, and that the detail is present in your visual experience. However, fixating on the period at the end of this sentence frustrates your ability to determine more than just a few surrounding words on the page. Furthermore, recent work on *change blindness* demonstrates that frequently we fail to notice prominent changes to a visual scene, such as a sailboat disappearing from an image, the replacement of a person with whom we are conversing, or the swapping of faces in a photograph (see Simons and Ambinder 2005; Simons and Rensink 2005). Moreover, aspects of a scene to which we do not attend, including those as striking as a gorilla strolling across the court during a basketball game, escape our notice, a phenomenon known as *inattentional blindness* (see Mack and Rock 1998; Simons and Chabris 1999; Chun and Marois 2002; Most et al. 2005). Noë (2004) has on these grounds criticized what he dubs the 'snapshot conception' of visual experience according to which one enjoys uniformly detailed visual awareness of an entire scene at any given moment. Others have argued for similar reasons that rich representations are absent from vision altogether (O'Regan 1992; O'Regan and Noë 2001). The evidence, however, is consistent with our registering the relevant information at a subpersonal (or even conscious) level but failing to retain, attend to, or access it (Simons and Rensink 2005; Prinz 2006). Even granting that we neither visually experience nor at any level visually represent informationally rich detail, one might revise one's characterization of representations to include mere sparse or incomplete detail (Lycan 2006). Indeed, this eases the computational and the explanatory burdens.

A second type of concern is that the received model leaves out the contributions of some factor critical to perceiving, such as movement, action, or the body. Thus, the received model is charged with failing to appreciate the dynamic quality of seeing as it unfolds over time and in response to a creature's engagement with its environment. The charge is that received theories fail because they consider vision merely as a static or disembodied phenomenon.

Such faults, however, do not belong to the computational framework itself. Amending an algorithm to incorporate relevant contributions and constraints might repair the defect. Ongoing research attempts to discover just such contributions. Even so, it is not clear that current models entirely fail to consider such contributions. For instance, distinguishing global patterns of optic flow that result from movement of the eyes (saccadic and intentional) and head from those that correspond to relative motion among objects is critical to determining the size, shape, and movement of objects in a scene. Capturing invariant environmental features by distinguishing patterns of stimulation that stem from changes to the environment from those due to a subject's activity is part of the task of perception on contemporary models. It is difficult to see how the proposed solutions fail to consider either the dynamic or action-involving character of that task.

Some might still object that information-processing accounts make perceiving an entirely subpersonal process, instead of an activity carried out by the creature itself. But subpersonal processes might be constitutive of perceiving, or might underlie it, without being identical with or providing the essence of perceiving. The representational view is compatible with understanding perception to involve the level of the person or creature, while particular informational theories of perceiving aim to explain how -- the mechanisms by which -- representing one's environment is possible given our physiology and more basic capacities.

A closely related worry about the received view is that it construes perceiving as something that depends entirety upon what takes place "in the head". Though characterizing or individuating perceptual content -- what is represented -- might require invoking external relations, perceptual experience on this view itself depends or supervenes entirely upon one's physiology (Lycan 2006). That is, stimulating the sensory surfaces, or fixing subsequent brain activity, suffices to generate a given perceptual experience. Critics, however, insist that such brain activity cannot be understood as constitutive of perceptual experience when considered in isolation from how it is embedded in a creature and an environment. The 'vehicles' of perceptual experience thus extend beyond the skin.

A growing cadre of philosophers, psychologists, and neuroscientists take the sum of these worries to constitute a strong case against the received view. They contend that no theory that frames perception in terms of sensory surface stimulation and subsequent rule-driven processing of internal representations captures what is distinctive about perceiving. They instead suggest that an adequate understanding of perception requires appreciating how a creature *uses* its body and its senses to interact with its environment. Deciphering perception thus requires comprehending "the level of detail of the biological machine" -- the level of *embodiment* (Ballard 1996: 461). The details of biological or physiological implementation are relevant on this view to understanding the tasks and algorithms for perceiving, contrary to Marr's strongly top-down theoretical approach.

According to one such theory, perceiving is a dynamic, purpose-driven way of interacting with the world. O'Regan and Noë's (2001) *sensorimotor account* holds that seeing is an activity that consists in a creature's exploring its environment in a skillful manner. They argue that although brain activity is a necessary part of what enables perception, no internal representation suffices for seeing. Seeing, for sensorimotor theorists, is not essentially mediated by detailed internal representations constructed "at some specific stage of neural processing" (Noë and Thompson 2002*b*: 6); it is essentially mediated only by one's implicit grasp upon *sensorimotor contingencies*, which are the ways that sensory stimulation varies in response to actions and movements. Contentful seeing is performance that exemplifies one's mastery of certain patterns of sensorimotor contingency. Perceiving an object requires exercising one's implicit understanding of the ways things look and feel from a variety of perspectives. Thus, one does not internally represent, but *enacts* perceptual content through performance. Perceiving, according to such an *enactive* conception, is a skill-based way of coming into contact with one's environment (see also Noë and Thompson 2002*a*; Noë 2004).

What we do, of course, makes a difference to what we see: turning my head leads to my seeing a window. Acting thus causally impacts what we perceive. But what we do may also be relevant to explaining *how* we see. At minimum, explaining vision requires positing perceptual principles that concern how sensory stimulation changes in response to what one does. Distinguishing subject-induced patterns of stimulation or optic flow from object-induced ones helps to ground visual awareness of a scene. So, proponents of the received view might reasonably grant that implicitly grasping, or subpersonally representing, sensorimotor contingencies is necessary for perceiving. The conceptual rift between perceiving and acting, since it is bridged by principles of vision, is therefore less sharp than early inquiry supposed. Still, it is a stretch to say, on such a conception, that action itself is constitutive of perception. For the received view, the boundary between perceptual processes and action remains intact because the principles in question are internal to vision. In brief, the relationship between perception and action remains causal (see Prinz 2006). According to ardent dissenters from the received view, nothing short of subsuming perceiving to acting explains perceiving.

Two steps toward resolution nevertheless serve to sharpen the focus of the conflict. First, clarifying the structure, detail, and accessibility of visual representations in framing the informational task of vision helps to address worries concerning how much we see. Second, recognizing movement-involving or motor-based constraints in the solution of the information-processing problem alleviates pressure to capture perception's dynamic and enactive characteristics. Such conciliatory steps suggest fruitful avenues for future research, but by no means dissolve this foundational dispute.

The varieties and bounds of perception

Up to now, I have discussed insights into the nature of perception that emerge from considering the kinds of functions perceptual processes perform and the general problems perceptual mechanisms must solve. To more carefully specify perceptual functions and to characterize its mechanisms, however, requires a scientific study of the sorts of things creatures perceive. At a lower level of abstraction, therefore, cognitive science sheds light on what it is to perceive by demanding and providing a systematic, scientific *accounting* of what we are capable of perceiving. Experimental psychology and philosophy increasingly attempt to answer such questions such as: What kinds of particulars and properties can we perceive? How much detail and variety is evident in perception?

What can humans perceive? One way to approach this question is to determine what we must mention in order to capture the accuracy conditions for perceptual states and experiences. Answering it is critical for articulating what capacities must be explained by a scientific theory of perception. It constrains what sorts of mechanisms must operate to enable perception, and it illuminates the boundaries of perceptual processes that must be respected by any account. Answers might vary from the austere to the liberal. If, for instance, what we perceive is relatively sparse, then a greater share of the work in acquiring knowledge about the world is left to post-perceptual inferences or other ampliative cognitive processes. The less we perceive, the more we must extrapolate or judge on the basis of relatively impoverished experience. If, on the other hand, perceptual experience is abundantly populated with qualities, features, and individuals, then the mechanisms involved must be appropriately sophisticated. For instance, it is one thing to say that we perceive qualities such as color and pitch, and quite another to say that we perceive the individual objects that bear them and about which we come to form beliefs. Perception's reach may even extend beyond ordinary material objects and their features: consider sounds, rainbows, shadows, and tastes.

Why does it matter how much we perceive and how much we infer? The answer has consequences for the *epistemic* relationship between perception and judgments or beliefs. The more we perceive, the more perception evidentially constrains judgment and belief. We might frame the issue as concerning the degree to which perception is "intelligent", or the degree to which it matches, and thus compels, the content of thought. Suppose you judge that a tomato is present before you, or that it is roundish, or that a causal collision takes place when it hits the ground, on the basis of perceptually experiencing that object to be before you, or to have those features, or to take part in a collision. Those judgments differ in evidential status from equivalent ones made on the basis of experiencing just a shifting pattern of light and colors over time. That is, *perceptually experiencing* a red tomato to collide with the floor grounds and constrains your forming the belief that the tomato hit the floor in a way that merely seeing fleeting patterns of colors does not. Forming the belief on the basis of the latter experience requires the deployment of inference and thought. Forming the belief when you experience the tomato crash requires only accepting what you see, and inherits its reliability from that perceptual experience. One significant upshot of the difference is that the sparse experience, but not the rich one, is compatible with a wide range of beliefs about the environment, such as that there is an image, a hologram, or a wax fruit ahead. The rich experience is accurate only if a tomato lies ahead. Where perceptual capacities end and higher-level cognition begins therefore matters to how we envision the intelligence of perception. In principle, investigating these issues through empirical research helps to resolve disputes among theories of perception and perceptual content. One prominent methodology involves constructing and manipulating perceptual illusions. Another involves setting tasks that require the resources and response times of perception. Such illusions and performance measures, which are to a great extent beyond a subject's control, provide evidence about the varieties and richness of perceptual experience that is accessible to researchers and subjects alike. Their interpretation and explanation offer clues to perception's functions and mechanisms.

Consider vision. What do we see? One might claim, as did early theorists, that we see only a pointillistic array of colors arranged to constitute an image-like visual field. A very austere account stops here, and leaves anything further to explicit recognition, judgment, or inference. Perception, such a view maintains, is a matter of being aware of qualitative features, while representing particulars such as objects and events bears the distinctive marks of conception or thought. Spelke (1988: 229) claims, "The parsing of the world into things may point to the essence of thought and to its essential distinction from perception." Recent work on the perception of surfaces (Nakayama et al. 1995) and objects (see Scholl 2002) suggests, however, that features alone do not suffice to capture the structure of visual perceptual experience, which appears to track individuals or objects according to Spelke's own principles of *contact*, *cohesion*, and spatio-temporal *continuity* (see Spelke 1990, 1994; Scholl and Leslie 1999). Pylyshyn (e.g., 2000, 2001, 2007) has proposed that explaining our ability to track multiple objects through time and feature changes requires positing *object indexes* fixed to up to four or five individuals at a time without representing their properties. Subsequent processes, such as selective attention, might then associate, or *bind*, feature information with visual indexes and serve as the basis for intermediate-level *object files*, or object representations that persist through feature changes and mediate recognition (see Leslie et al. 1998). Perceptual mechanisms found at the earliest stages of vision may carve the visual world in to *units* or *individuals*.

Even audition, the study of which until recently stopped at explaining awareness of audible qualities such pitch, timbre, and loudness, requires organization into perceptual units or "objects". Bregman (1990) argues that auditory perception involves carving or segmenting an auditory scene into discrete, temporally extended *streams*, which might be understood as individuals that bear audible qualities, have spatial locations, and persist through time. *Auditory scene analysis* is the process of extracting, from complex wave information arriving at the ears, information concerning the significant sounds and acoustic sources in one's environment. Auditory scene analysis, for instance, grounds one's ability to hear two simultaneous sounds as distinct. Perceptual "units" therefore encompass categories beyond just visual or ordinary objects, and include such individuals as sounds and events (see Gaver 1993; Kubovy and Van Valkenburg 2001; Griffiths and Warren 2004; O'Callaghan 2007). Far from being the result of post-perceptual processes, such units serve to structure perceptual processes and experience.

Recent work on a surprising class of *cross-modal* perceptual illusions suggests that perceptual units need not even be unique to a given perceptual modality. Consider the well-known ventriloquist illusion, in which seeing the location of an apparent sound source illusorily alters the experienced location of a sound (Bertelson 1999); the previously-mentioned McGurk effect, in which seeing a speaker move her lips alters the phoneme one hears (McGurk and MacDonald 1976); or the striking, recently discovered sound-induced flash illusion, in which hearing two beeps accompanied by a single flash leads one illusorily to visually experience two flashes (Shams et al. 2000, 2002). Each involves sense perception that occurs through one modality impacting experience ordinarily associated with another modality. Explaining the intelligibility of such illusions may require attributing to systems of sense perception, considered collectively, a multi-modal or modality-independent grasp upon common sources of sensory stimulation across modalities (Welch and Warren 1980; Bertelson and de Gelder 2004; O'Callaghan 2007, ch. 11, forthcoming).

In light of its history of conservatism, what is perhaps most remarkable about current research into perception is the arsenal of *perceptual capacities* that researchers have uncovered. To take a prominent recent example, Scholl and colleagues have argued, following Michotte (1963), that *causation* is among the things we perceive (see, e.g., Scholl and Tremoulet 2000). Not only do causal relations factor into intuitive reports concerning the content of perceptual experience, but causal interactions generate certain perceptual illusions. For instance, visually experiencing an apparent collision, in which two objects do not touch but exhibit motion characteristic of rebounding, leads subjects to underestimate the distance between the objects at the time of the apparent interaction (Scholl and Nakayama 2004). Furthermore, perceptual tracking of objects is influenced by apparent causal interactions. Take two objects moving toward each other from either side of a screen. Hearing a beep at the time when the objects meet prompts subjects to visually experience the objects to collide and rebound, instead of to continue past each other on their respective paths (Sekuler et al. 1997; Watanabe and Shimojo 2001). Even philosophers, counter to longstanding skepticism stemming from Hume about our capacity to perceive causal interactions, recently have argued that visual experience sometimes represents causation (Siegel forthcoming). Experiencing causal relations, it appears, is effortless, automatic, and irresistible.

The role of phenomenology

As the foregoing makes evident, philosophers and psychologists frequently rely on conscious or introspectible aspects of experience to launch or constrain theoretical discussion. Introspection and reliance upon perceptual phenomenology, however, have come under attack. Though the phenomenology of a perceptual experience is supposed to concern what it is like for a subject to have that experience, phenomenological reports often are incomplete, inaccurate, or unreliable (see Jack and Roepstorff 2003; Roepstorff and Jack 2004). Though we take ourselves to have detailed awareness of our surroundings at any given moment, and though we believe that we know just what we see, careful investigation challenges both of these beliefs. Some researchers have suggested that the sense of a vivid phemenology of visual experience is a "grand illusion" (see Noë 2002). How could any methodology that depends upon elusive phenomenology and unreliable introspection be trusted to play such a central role?

To be clear, nonveridical experiences such as illusions show that perceptual phenomenology does not suffice for perceiving. A more challenging question is whether phenomenology is necessary for perceiving. Perceptual experiences do frequently accompany perception, but are they required? Recent evidence suggests otherwise for at least some instances. Certain subjects with primary visual cortex damage respond to their environment in ways characteristic of perceiving without introspectible conscious experience. Such *blindsighted* subjects reliably form beliefs about things before their eyes without reporting seeing them (see, e.g., Weiskrantz 1986; Block 1995; Boyer et al. 2005). Furthermore, patients with a form of visual agnosia report no awareness of certain spatial features, yet those features appear successfully to guide their action (Milner and Goodale 1995). Even ordinary perceivers subject to certain visual spatial illusions act in ways appropriate to the scene's true geometry, such as by adjusting grasp width to the actual size of Titchener circles (Goodale and Humphrey 1998). Subjects sometimes make judgments and act on the basis of sensory information, though they lack corresponding conscious awareness. It would be dogmatic to deny that such subjects perceive, in some fashion.

If phenomenological reports are unreliable, and if perceptual phenomenology is neither sufficient nor necessary for perceiving, what is the role of phenomenology and first-person experience in the philosophical and empirical investigation of perception and its contents? As a start, it is plausible that differences to phenomenology, or differences to how things seem, imply differences to the content of perceptual experience or to how things perceptually are represented to be (Byrne 2001). This observation by itself need not imply infallible introspective access to the respective phenomenological features or contents. Rather, we must only be capable of detecting contrasts. One reasonable strategy involves pointing to an introspectible phenomenological contrast and suggesting that a given representational difference best explains that contrast (see Siegel 2006*a*, *b*, forthcoming). This requires from phenomenology only that apparent phenomenological differences imply genuine phenomenological contrasts. The explanation for the contrast must stand on further theoretical, and not mere phenomenological, grounds.

Consulting phenomenology also can be an important tool for investigating perception's mechanisms. Nakayama et al. (1995: 21-2) argue that phenomenologically grounded results are more objective than many suppose and suggest three reasons phenomenological methods are valuable. First, surprisingly wide agreement with respect to phenomenology occurs with well-crafted experiments and demonstrations. Numerous demonstrations evoke near universal agreement, in contrast to many other psychological methods, which rely upon subtle statistical analysis of large data sets. Second, compelling phenomenological results are immediate and verifiable across subjects. Researchers and audiences can view results for themselves. Finally, such methods provide a large, diverse database of results with relative ease and little cost. Overall, the results and conclusions drawn from phenomenological methods are widely circulated and accepted.

This alone does not justify their use to discover the deep mechanisms responsible for perceiving. For one thing, the structure of experience need not match that of perceptual processes. Such methods nonetheless furnish one form of data that must somehow be explained. Contrary to traditional views, however, the phenomenology of experience is often far from immediately obvious. Doing careful phenomenology takes considerable work. Responses based on reflection upon phenomenology should be treated as a kind of performance that might be attributed to a variety of factors apart from accurate reporting whose content matches the experience. If reports might be infused with information from other sources, such as one's background beliefs concerning the items in a scene, or some strategy adopted to respond to ambiguous experiences, then perhaps no unique, epistemically privileged level of introspectively accessible phenomenology exists. It is unwise in both philosophy and in science to rely exclusively on phenomenological reports and reflections -- those of others or one's own. It is, however, compelling to understand them as data that inform the construction of philosophical and scientific theories of perception. It remains, after all, a goal of investigating perception to explain the seemings.

References

Ballard, D. (1996). On the function of visual representation. In Akins, K., editor, *Perception*, volume 5 of *Vancouver Studies in Cognitive Science*, pages 111–131. Oxford University Press, Oxford. Reprinted in Noë and Thompson (2002*a*).

Bertelson, P. (1999). Ventriloquism: A case of cross-modal perceptual grouping. In Aschersleben, G., Bachmann, T., and Müsseler, J., editors, *Cognitive Contributions to the Perception of Spatial and Temporal Events*, pages 347–362. Elsevier, Amsterdam.

Bertelson, P. and de Gelder, B. (2004). The psychology of multimodal perception. In

Spence, C. and Driver, J., editors, *Crossmodal Space and Crossmodal Attention*, pages 141–177. Oxford University Press.

Block, N. (1990). Inverted earth. In Tomberlin, J., editor, *Philosophical Perspectives*, volume 4. Ridgeview, Atascadero, CA.

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18:227–247.

Block, N. (1996). Mental paint and mental latex. *Philosophical Issues*, 7:19–49.

Block, N. (forthcoming). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*.

Boyer, J. L., Harrison, S., and Ro, T. (2005). Unconscious processing of orientation and color without primary visual cortex. *Proceedings of the National Academy of Science*, 102(46):1675–1679.

Bregman, A. S. (1990). *Auditory Scene Analysis: The Perceptual Organization of Sound*. MIT Press, Cambridge, MA.

Byrne, A. (2001). Intentionalism defended. *Philosophical Review*, 110:199–240.

Chalmers, D. J. (2004). The representational character of experience. In Leiter, B., editor, *The Future for Philosophy*, pages 153–181. Oxford University Press, Oxford.

Chalmers, D. J. (2006). Perception and the fall from Eden. In Gendler, T. S. and Hawthorne, J., editors, *Perceptual Experience*, pages 49–125. Clarendon Press, Oxford.

Chun, M. M. and Marois, R. (2002). The dark side of visual attention. *Current Opinion in Neurobiology*, 12(2):184–189.

Dennett, D. C. (1987). *The Intentional Stance*. MIT Press, Cambridge, MA.

Dennett, D. C. (1991). *Consciousness Explained*. Back Bay Books, Boston.

Descartes, R. (1637/1988). Optics. In Cottingham, J., Stoothoff, R., and Muroch, D., editors, *Descartes: Selected Philosophical Writings*. Cambridge University Press, Cambridge, UK.

Dretske, F. (1995). *Naturalizing the Mind*. MIT Press, Cambridge, MA.

Fodor, J. (1975). *The Language of Thought*. Harvard University Press, Cambridge, MA.

Fodor, J. (1983). *The Modularity of Mind*. MIT Press, Cambridge, MA.

Fodor, J. A. and Pylyshyn, Z. W. (1981). How direct is visual perception? : Some reflections on Gibson's "ecological approach". *Cognition*, 9:139–196.

Foster, J. (2000). *The Nature of Perception*. Clarendon Press, Oxford.

Gaver, W. W. (1993). What in the world do we hear? : An ecological approach to auditory event perception. *Ecological Psychology*, 5(1):1–29.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Lawrence Erlbaum, Hillsdale, NJ.

Goodale, M. A. and Humphrey, G. K. (1998). The objects of action and perception. *Cognition*, 67:181–207.

Griffiths, T. D. and Warren, J. D. (2004). What is an auditory object? *Nature Reviews Neuroscience*, 5:887–892.

Harman, G. (1990). The intrinsic quality of experience. In Tomberlin, J., editor, *Philosophical Perspectives*, volume 4. Ridgeview, Atascadero, CA.

Hatfield, G. (2002). Perception as unconscious inference. In Heyer, D. and Mausfeld, R., editors, *Perception and the Physical World: Psychological and Philosophical Issues in Perception*. John Wiley & Sons, Ltd.

Helmholtz, H. L. (1967/1910). *Hadnbuch der physiologischen Optik*. L. Voss, Leipzig.

Hinton, J. M. (1973). *Experiences*. Clarendon Press, Oxford.

Jack, A. I. and Roepstorff, A. (2003). Why trust the subject? *Journal Of Consciousness Studies*, 10(9-10):V–XX.

Jackson, F. (1977). *Perception: A Representative Theory*. Cambridge University Press, Cambridge.

Johnston, M. (2006). Better than mere knowledge: the function of sensory awareness. In Gendler, T. S. and Hawthorne, J., editors, *Perceptual Experience*. Clarendon Press, Oxford.

Kubovy, M. and Van Valkenburg, D. (2001). Auditory and visual objects. *Cognition*, 80:97–126.

Leslie, A. M., Xu, F., Tremoulet, P. D., and Scholl, B. J. (1998). Indexing and the object concept: developing 'what' and 'where' systems. *Trends in Cognitive Sciences*, 2(1):10–18.

Scholl, B. J. and Leslie, A. M. (1999). Explaining the infant's object concept: Beyond the perception/cognition dichotomy. In Lepore, E. and Pylyshyn, Z., editors, *What is cognitive science?*, pages 26–73. Blackwell, Oxford.

Lewis, D. (1966). Percepts and color mosaics in visual experience. *The Philosophical Review*, 75:357–368.

Loar, B. (2003). Transparent experience and the availability of qualia. In Smith, Q. and Jokic, A., editors, *Consciousness*. Oxford University Press, Oxford.

Lycan, W. (1996). *Consciousness and Experience*. MIT Press, Cambridge, MA.

Lycan, W. (2006). Enactive intentionality. *PSYCHE*, 12(3).

Mack, A. and Rock, I. (1998). *Inattentional Blindness*. MIT Press, Cambridge, MA.

Marr, D. (1982). *Vision*. W. H. Freeman, San Francisco, CA.

Marr, D. and Nishihara, H. K. (1978). epresentation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London, Series B, Biological Sciences*, 200(1140):269–294.

Martin, M. G. F. (2002). The transparency of experience. *Mind and Language*, (17):376–425.

Martin, M. G. F. (2004). The limits of self-awareness. *Philosophical Studies*, 120:37–89.

McDowell, J. (1982). Criteria, defeasibility, and knowledge. *Proceedings of the British Academy*, 68:455–79.

McDowell, J. (1994). The content of perceptual experience. *Philosophical Quarterly*, 44:190–205.

McGurk, H. and MacDonald, J. (1976). Hearing lips and seeing voices. *Nature*, 264:746–748.

Michotte, A. (1963). *The Perception of Causality*. Basic Books, New York.

Milner, A. D. and Goodale, M. A. (1995). *The Visual Brain in Action*. Oxford University Press, Oxford.

Most, S. B., Scholl, B. J., Clifford, E. R., and Simons, D. J. (2005). What you see is what you set: sustained inattentional blindness and the capture of awareness. *Psychological Review*, 112(1):217–242.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83(4):435–450.

Nakayama, K., He, Z. J., and Shimojo, S. (1995). Visual surface representation. In Kosslyn, S. M. and Osherson, D. N., editors, *Visual Cognition*, volume 2 of *An Invitation to Cognitive Science*. MIT Press, second edition.

Noë, A. (2002). Is the visual world a grand illusion? *Journal of Consciousness Studies*, 9(5/6):1–12.

Noë, A. (2004). *Action in Perception.* MIT Press, Cambridge, MA.

Noë, A. and Thompson, E., editors (2002*a*). *Vision and Mind: Selected Readings in the Philosophy of Perception.* MIT Press, Cambridge, Massachusetts.

Noë, A. and Thompson, E. (2002*b*). Introduction. In Noë and Thompson (2002*a*).

O'Callaghan, C. (2007). *Sounds: A Philosophical Theory*. Oxford University Press, Oxford.

O'Callaghan, C. (forthcoming 2008). Seeing what you hear: crossmodal illusions and perception. *Philosophical Issues*.

O'Regan, J. K. (1992). Solving the 'real' mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46:461–488.

O'Regan, J. K. and Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24:939–1031.

Peacocke, C. (1983). *Sense and Content*. Oxford University Press, Oxford.

Prinz, J. (2006). Putting the brakes on enactive perception. *PSYCHE*, 12(1).

Pylyshyn, Z. W. (1984). *Computation and Cognition: Toward a Foundation for Cognitive Science*. MIT Press, Cambridge, MA.

Pylyshyn, Z. W. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22(3):341–423.

Pylyshyn, Z. W. (2000). Situating vision in the world. *Trends in Cognitive Sciences*, 4(5):197–207.

Pylyshyn, Z. W. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition*, 80:127–158.

Pylyshyn, Z. W. (2007). *Things and Places: How the Mind Connects With the World.*

Jean Nicod Lecture Series. MIT Press, Cambridge, MA.

Robinson, H. (1994). *Perception*. Routledge, London.

Rock, I. (1983). *The Logic of Perception*. The MIT Press, Cambridge, MA.

Roepstorff, A. and Jack, A. I. (2004). Trust or interaction? *Journal Of Consciousness Studies*, 11(7-8):V–XXII.

Scholl, B. and Nakayama, K. (2004). Illusory causal crescents: Misperceived spatial relations due to perceived causality. *Perception*, 33:455–469.

Scholl, B. and Tremoulet, P. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences*, 4(8):299–309.

Scholl, B. J., editor (2002). *Objects and Attention*. MIT Press, Cambridge, MA.

Sekuler, R., Sekuler, A. B., and Lau, R. (1997). Sound alters visual motion perception. *Nature*, 385:308.

Shams, L., Kamitani, Y., and Shimojo, S. (2000). What you see is what you hear. *Nature*, 408:788.

Shams, L., Kamitani, Y., and Shimojo, S. (2002). Visual illusion induced by sound. *Cognitive Brain Research*, 14:147–152.

Siegel, S. (2005). The contents of perception. In Zalta, E. N., editor, *Stanford Encyclopedia of Philosophy*.

Siegel, S. (2006*a*). Subject and object in the contents of visual experience. *Philosophical Review*, 115(3):355–388.

Siegel, S. (2006*b*). Which properties are represented in perception? In Gendler, T. S. and Hawthorne, J., editors, *Percepual Experience*. Oxford University Press, New York.

Siegel, S. (forthcoming 2008). The visual experience of causation. *Philosophical Quarterly*.

Simons, D. J. and Ambinder, M. S. (2005). Change blindness: Theory and consequences. *Current Directions in Psychological Science*, 14(1):44–48.

Simons, D. J. and Chabris, C. F. (1999). Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception*, 28:1059–1074.

Simons, D. J. and Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences*, 9(1):16–20.

Spelke, E. S. (1988). Where perceiving ends and cognition begins: the apprehension of objects in infancy. In Yonas, A., editor, *Perceptual Development in Infancy*, pages 197–234. Lawrence Erlbaum, Hillsdale, NJ.

Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, 14:29–56.

Spelke, E. S. (1994). Initial knowledge: Six suggestions. *Cognition*, 50:431–445.

Strawson, P. F. (1979). Perception and its objects. In MacDonald, G. F., editor, *Perception and Identity: Essays Presented to A. J. Ayer with his Replies*. Cornell University Press, Ithaca.

Tye, M. (1995). *Ten Problems of Consciousness*. MIT Press, Cambridge, MA.

Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press, Cambridge, MA.

Tye, M. (2003). Qualia. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*.

Watanabe, K. and Shimojo, S. (2001). When sound affects vision: effects of auditory grouping on visual motion perception. *Psychological Science*, 12:109–116.

Weiskrantz, L. (1986). *Blindsight: A Case Study and Implications*. Clarendon Press, Oxford.

Welch, R. B. and Warren, D. H. (1980). Immediate perceptual response to intersensory discrepancy. *Psychological Bulletin*, 88(3):638–667.