PERCEPTION
Cambridge Handbook to Cognitive Science
Eds. Keith Frankish and William Ramsey
Casey O'Callaghan

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 cannot tell definitively if you are hallucinating, how are you ever acquainted with things like tomatoes, barns, collisions, colors, sounds, and odors?

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, for instance, do your body and brain, which are made of cells and neural tissue, ground your awareness of the sights and sounds of a tomato being squashed?

This chapter introduces the traditional philosophical problem of perception, which concerns whether our naïve 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 it provides the conceptual backdrop to contemporary discussions among philosophers and cognitive scientists about what it is to perceive. The chapter next turns to empirically motivated theoretical issues about perception. These concern how to characterize the tasks of perception; alternative ways to understand the role of processes and mechanisms involved in perceiving; and the relationships among perception, other forms of cognition, and action. In light of methodological concerns, the chapter concludes by discussing 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, audio of the syllable /ba/ sounds like /da/ when viewing a speaker's mouth pronounce /ga/. Such illusions involve perceiving a thing while misperceiving some of its features. Hallucinations, however, involve complete figments of experience, as when Macbeth 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 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 ordinary things and qualities (like tomatoes, colors, and sounds) are not the objects of illusory and hallucinatory perceptual experiences, and if nothing subjectively accessible indicates 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 illusion and hallucination that is undetectable by the subject. It is useful to classify them according to whether the objects of perception, if any, are internal or external (that is, whether they are mind-dependent or mind-independent), and according to whether perceptual awareness of things in the world, if any, is mediated awareness or not (that is, 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 (see, e.g., Smith 2002; Crane 2008).

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 things like tables and chairs. Instead, they are experience-dependent items of sense that are present to one equally in veridical and delusive experience. 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. Perception thus is mediated by awareness of internal features present in sensation (see, e.g., Jackson 1977; Foster 2000).

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. Perceptual experiences require a certain kind of relation between a subject and an ordinary subject-independent object, in which the object is partly constitutive of the perceptual experience. Delusions do not involve such a constitutive relation to an ordinary subject-independent object. Perception on this account is direct because it requires no awareness of internal mediators. According to disjunctive theories, however, subjects are at best in a position to characterize what they undergo as, for example, *indistinguishable from when I see a*

tomato smash—one either hallucinates or experiences a tomato (see Haddock and Macpherson 2008; Byrne and Logue 2009).

Intentional, or *representational*, 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 veridical experiences and subjectively indistinguishable illusions belong to some common explanatory psychological kind. According to intentionalism, perceiving, like believing and desiring, involves representing things to be a certain way. 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, thus 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. According to the intentionalist, however, *perceptual experience*, like belief, 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 explaining why 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 (see Harman 1990; Tye 2000).

If 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. In addition, perceptual experiences count as sources and reasons for judgments and beliefs, and bear distinctive evidential weight. Experiences, however, need not compel belief or judgment. One might withhold commitment to the truth of what one experiences, and perhaps one need not even acquire a disposition to judge or so commit. 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 bodily sensations, such as those you are aware of when you experience pains, nausea, or dizziness, seem private or not independent from oneself. First, they do not seem to be

located outside one's body. Second, their existence seems to depend on their being experienced: a pain exists if one hurts, and unfelt dizziness is no dizziness at all. Third, more controversially, their experience lacks accuracy conditions: one cannot misperceive a pain, and illusory nausea is just plain nausea.

Some philosophers hold that sense perception involves consciousness of purely qualitative features, or *qualia*, that are intrinsic to one's experience (see Jackson 1982). Though it is a matter of controversy whether any such features exist, they are at least quite difficult to discern *as such*. The fact that one "sees right through" one's experiences to the world has been dubbed the *diaphanousness* or *transparency* of perceptual experience. As a result, one's attempts to introspectively attend to features intrinsic to a perceptual experience nearly always deliver either features of things one seems to perceive, such as colors and shapes, or extrinsic, relational features of one's experience, such as that one experiences a red thing. According to intentionalists about the phenomenological character of perceptual experience, any qualia we do sense strike us either as qualities of things we perceive, or as representational features of experiences (see, e.g., Harman 1990; Tye 2003; cf. Loar 2003).

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 neural 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. 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 micro-saccades that occur up to 60 times per second) the image moves continuously relative to the retina. Rod and cone receptors, which are sensitive to different wavelengths of light, are distributed unevenly. Image information is lost where the optic nerve departs the retina, though we don't experience a "blind spot" in our vision. 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 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 perceiving. On the other hand, we might think that although these mechanisms reveal a critical part of what enables perception, citing internal processes as such misses some essential aspect of what it is to perceive. As the following three sections mean to make clear, 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.

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 conscious or aware. But even consciously sensing an image requires extracting features such as color, shape, and motion from light. Illumination upon the retina, for instance, is a product of ambient light and surface reflectance, and thus underdetermines color experience. Even if a scene's colors were simply projected onto a retinal image, being conscious of them would require an additional mind's eye. 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. As above, perceiving requires something akin to problem solving because the information on which it is based, even if that is an image, underdetermines what one perceives. Early theorists took perceiving, understood as discerning or grasping an image's significance, to deploy general-purpose cognitive capacities that are not peculiar to perceiving, but which are involved in many varieties of thinking and reasoning. 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 or from sensory stimulation. One might suggest that perceptions, concerning, for instance, size and distance, result from mechanisms such as learning, association, inference, or intuition. But it is 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, or if the only information available is how a sense organ is stimulated.

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, according to Helmholtz, such inferences are based upon sensations that possess phenomenal features, and since Helmholtz relies on a general, associative model of inference, his view exhibits the failings characteristic of early theories. 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. It prefigures contemporary accounts according to which *subpersonal* sensory processes fuel unconscious transitions akin to a form of deductive or inductive inference (see Rock 1983).

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, they hold that specialized subsystems, and not a general rational or cognitive faculty, conduct perceptual processing. Deploying general-purpose problem-solving strategies would be radically inefficient given the specificity and complexity of perceptual tasks. To efficiently transform light information into full-fledged perceptions of public objects and features requires extracting, representing, and putting to use specific kinds of information at different stages, and it requires specialized rules and assumptions to guide the transitions. As a result, perceptual systems are, to a significant degree, unaffected by one's beliefs and reasoning. Though they furnish materials for thought and action, they are to a great extent *modular* or informationally encapsulated (Fodor 1983; Pylyshyn 1984). Perceptual capacities on this model forfeit the generality of those proposed by Helmholtz and other early theorists.

According to the contemporary establishment view, perceiving involves processes that resemble judgment and reasoning. The degree of specialization in perceptual systems, however, makes perception more automatic than classic views supposed. Whether 'inference' is understood literally 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. 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 involves the kind of information processing that is characteristic of representational systems, such as computers. Perceiving, on this account, is constructing a representation of one's environment from impoverished sensory 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. One might also wonder why *conclusions* of such inferences are identical with or lead to perceptual *experiences*.

Gibson and Marr: direct pickup and computation

The view that perceiving 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 a matter of representing, transforming, and augmenting impoverished information drawn from the senses. There is no need, Gibson claims, to infer or otherwise intelligently construct a rich internal description or representation of the world. The world, available directly to be perceived, eliminates the need for representations; it serves as its own model.

Gibson suggests that unconscious inferences are unnecessary for vision since information concerning features that matter to the creature is present in the pattern of light that reaches the eye, or the *ambient optical array*. The key is that 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 the objects and features in the environment. For instance, changes in illumination cause relatively uniform changes across the optic array. And the pattern of optic flow when an object moves differs from when the creature moves. When an object moves, it produces a local, relative change to an otherwise static optic array. When the

creature moves, it produces a distinctive global pattern of optic flow. Walking forward, for instance, creates flow outward from a vanishing point in the direction of travel. 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. It is not a matter of constructing representations by means of unconscious inferences from impoverished, static sensory information. According to Gibson, subpersonal processes *enable* perceiving. However, for Gibson, perceiving is an achievement of nothing less than the creature in its environment.

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 information 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). So Marr 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 also 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 program) addresses the format in which to represent each sort 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 that reflect one's environment. Assumptions about the natural world—concerning, for example, how scene geometry projects to the retinal image (for instance, how a three-dimensional object projects to a two-dimensional array) or that rigid rather than flexible bodies produce a given pattern of

stimulation over time—are encoded by the visual system and 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 increasingly fine 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 one predominant contemporary approach to perception in cognitive science.

Representing and enacting

According to the received approach, perceiving is tantamount to representing. It principally involves constructing a representation of the immediate environment from the noisy, ambiguous stimulation of sensory surfaces. According to this conception, all

perceptual awareness is, in a sense that may not be phenomenologically accessible to the subject, mediated by representations derived subpersonally 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 right now that you are seeing all of the words on this page, and that the details are all present in your visual experience. However, fixating on the period at the end of this sentence frustrates your attempts to recognize 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, e.g., 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* (Mack and Rock 1998; Simons and Chabris 1999). On these grounds, Noë (2004) criticizes 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. Some have argued for similar reasons that rich representations are absent from vision altogether (e.g., 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. Even granting that we neither visually experience nor, at any level, visually represent in information-rich detail, one might revise one's characterization of representations to include mere sparse or incomplete detail. 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, for instance, the received model has been charged with failing to appreciate the dynamic quality of seeing as it unfolds over time or in response to a creature's engagement with its environment. The charge is that establishment theories fail because they consider vision merely as a *static* or *disembodied* phenomenon. Establishment theorists, however, might respond as follows. Such faults do not belong to the overall computational framework itself, but rather to specific algorithms within that framework. Amending an algorithm to incorporate the relevant contributions and constraints might repair the defect. Ongoing research attempts to discover just such contributions. Moreover, it is not even clear that current models entirely fail to consider such contributions. For instance, distinguishing global patterns of optic flow that result from movements of the head and eyes (saccadic and intentional) from patterns that correspond to relative motion among objects is critical to determining the size, shape, and movement of objects in a scene. So, on contemporary models, detecting invariant features of the environment by distinguishing patterns of stimulation caused by changes to the environment from those caused by a subject's activity is part of the task of perception. Thus, the proposed solutions might after all

incorporate resources that help explain the dynamic and action-involving character of perception.

Some might still object that the 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—by explicating the mechanisms by which—representing one's environment is possible given our more basic capacities and our physiology.

A 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 the received view itself depends or supervenes entirely upon one's physiology. Stimulating sensory surfaces causes brain activity; such brain activity grounds perceptual processing that suffices to generate a given perceptual experience. Critics, however, insist that neither brain activity nor internally-grounded processing alone, considered in isolation from how it is embedded in a creature and an environment, constitutes perceptual experience. The *vehicles* of perceptual experience thus extend beyond the brain, or even beyond the skin.

A growing cadre of philosophers, psychologists, and neuroscientists takes 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. Instead, an adequate understanding of perception requires appreciating how a creature *uses* its body and its senses to interact with its environment. There are two broad themes in this work. One is that deciphering perception requires comprehending "the level of detail of the biological machine" (Ballard 1996: 461)—how a creature is *embodied*. Another is that it requires appreciating the way in which a creature is *embedded* or *situated* in its environment.

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, for instance, is an activity that consists in a creature's exploring its environment in a skillful manner (see also Hurley 1998; Noë 2004). They argue that although brain activity is a necessary part of what enables perception, no internal representation suffices for seeing. Seeing, for the sensorimotor theorist, is not essentially mediated by detailed internal representations constructed "at some specific stage of neural processing" (Noë and Thompson 2002: 6). It is essentially mediated only by the exercise of one's implicit grasp of *sensorimotor contingencies*, which are the ways that sensory stimulation varies in response to actions and movements. Perceiving thus exemplifies one's implicit understanding of the ways things look and feel from a variety of perspectives. One does not internally represent, but *enacts* perceptual content through skillful performance.

Perceiving, according to such an *enactive* conception, is a skill-based way of coming into *contact* with one's environment.

What we do, of course, makes a difference to what we see: turning my head leads to my seeing a jukebox. 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. Proponents of the establishment view may 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, therefore is less sharp than earlier 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 (see, e.g, Prinz 2006). In brief, the relationship between perception and action remains causal. According to ardent dissenters from the received view, perceiving must be subsumed under acting.

With this in mind, two steps toward resolution nevertheless serve to sharpen the focus of the conflict. The first step is to address worries concerning *how much* we see in framing the informational task of vision. In particular, clarify the amount of detail (rich or poor) in visual representations; specify the level (subpersonal or personal) of given representational states; and state whether such representations are accessible to other cognitive operations and to which ones, such as short-term recall or explicit reasoning, they are accessible. This helps address concerns about the snapshot conception of experience, change blindness, and inattentional blindness. The second step is to recognize movement-involving or motor-based constraints in the solution of the information-processing problem. For instance, seeing external objects depends on grasping how not just the environment but also a creature's movements impact optic flow. Seeing involves distinguishing creature-generated from environment-generated patterns of sensory stimulation. This alleviates some pressure to capture perception's dynamic and enactive characteristics. While such conciliatory steps suggest fruitful avenues for future research, they by no means dissolve this foundational dispute. The remaining dispute turns on at least three questions. First, is action constitutive of perception? Second, does what occurs in the head suffice for having perceptual experiences? Third, is perception a subject's direct contact with the world, or is it mediated by representations?

The role of phenomenology

This chapter has discussed philosophical and theoretical questions about the nature of perception that emerge from considering the role of perceptual illusions, the functions perception performs, and the general problems perceptual mechanisms solve. To more carefully characterize perceptual content, specify perceptual functions, and explicate its mechanisms requires a scientific study of the sorts of things creatures perceive. So, at a more concrete level, 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? And, how much detail and variety is evident in perception?

In each of these enterprises, philosophers and psychologists frequently rely on conscious or introspectible aspects of experience to launch or constrain theoretical discussion. What is the proper role of *phenomenology*, and what are the limits of appeals to phenomenology in theorizing about perception (see also Jack and Roepstorff 2003; Roepstorff and Jack 2004; Noë 2007)?

Non-veridical 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? Evidence suggests otherwise in at least some cases. 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 (Weiskrantz 1986). 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 rather than to their apparent size (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 so, perceptual phenomenology is neither sufficient nor necessary for perceiving.

Nakayama et al. (1995: 21-2) nevertheless argue that consulting phenomenology is an important tool for investigating perception. They claim that phenomenologically grounded results are more objective than many suppose and present three reasons phenomenological methods are valuable. First, well-crafted experiments and demonstrations lead to surprisingly wide agreement with respect to phenomenology. They note that 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. Nakayama et al. argue that phenomenological methods are widespread for good reason and suggest that the results and conclusions they ground are well founded.

However, reliance upon introspection and reports of perceptual phenomenology has come under attack (see, e.g., Schwitzgebel 2008). 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. Though many

of us take ourselves to have detailed awareness of our surroundings at any given moment, and believe we know just what we see, careful investigation challenges both of these beliefs. Some researchers suggest the impression that visual experience has a vivid, introspectible phemenology is a "grand illusion" (see Noë 2002). How could any methodology that depends upon elusive phenomenology and unreliable introspection be trusted?

As a start, it is plausible that differences in phenomenology, or differences in what it is like to have a perceptual experience, imply differences in the content of perceptual experience, or to how things perceptually are represented to be (Byrne 2001). This does not by itself imply infallible introspective access to determinate phenomenological features or to the contents of experiences. But, suppose we are capable under some conditions of detecting phenomenological *contrasts*. One strategy that uses phenomenology reasonably to get at perceptual content is to point to some introspectible phenomenological contrast and to argue that a given representational difference best explains that contrast (see Siegel 2006). This requires from phenomenology only that phenomenological or introspective reports do not wildly over-ascribe apparent phenomenological differences. Both the explanation for the contrast and the claim about perceptual content must stand on further theoretical, and not mere phenomenological, grounds.

This does not justify using phenomenology and first-person methods to discover the deep mechanisms responsible for perceiving. The structure of experience need not match that of perceptual processes. Such methods nonetheless furnish data that somehow must be explained. If an experience report is accurate, it provides evidence about the product of perceptual mechanisms. Contrary to traditional views associated with Descartes, however, the phenomenology of experience often is not immediately obvious. Doing careful experience reporting takes considerable work. Responses based on phenomenological reflection should be treated as a kind of performance that might be attributed to a variety of factors apart from accurately reporting perceptual experiences. 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—someone else's or one's own. It is, however, compelling to understand introspective reports 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.

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

Byrne, A. and Logue, H., editors (2009). *Disjunctivism: Contemporary Readings*. MIT Press, Cambridge, MA.

Crane, T. (2008). The problem of perception. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*.

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

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.

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.

Haddock, A. and Macpherson, F., editors (2008). *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press, Oxford.

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. (1867). *Handbuch der physiologischen Optik*. L. Voss, Leipzig.

Hurley, S. L. (1998). *Consciousness in Action*. Harvard University Press, Cambridge, MA.

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.

Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32(127): 127-136.

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.

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). Representation 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.

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

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., editor (2007). *Phenomenology and the Cognitive Sciences* (special issue on Dennett's heterophenomenology), 6: 1–270.

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

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.

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.

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.

Schwitzgebel, E. (2008). The unreliability of naive introspection. *Philosophical Review*, 117:245–273.

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

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

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.

Smith, A. D. (2002). *The Problem of Perception*. Harvard University Press, Cambridge, MA.

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

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

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

Further reading

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

Bregman's work brought techniques and theoretical resources familiar in contemporary cognitive science to bear on problems concerning auditory perception.

Gendler, T. S. and Hawthorne, J., editors, (2006). *Perceptual Experience*. Oxford University Press, New York.

A collection of new essays that provides a cross-section of recent philosophical work on the nature of perceptual content and perceptual experience.

Gibson, J. J. (1966). *The Senses Considered as Perceptual Systems*. Houghton Mifflin, Boston, MA.

One of the most influential books on sensory perception. It articulates the framework of Gibson's ecological approach to perception.

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

A recent interdisciplinary collection of essays on foundational and theoretical issues in the cognitive science of perception.

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

Argues from evidence that subjects fail to notice features of a scene to which they do not attend that conscious perceptual awareness requires attention.

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

Marr's book is the seminal presentation of the computational approach to visual representation, information processing, and perception.

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

A vivid, influential, and accessible recent introduction and defense of the enactive view of perception.

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

A selection of classic contemporary essays concerning philosophical issues about the nature of visual perception.

Palmer, S. (1999). *Vision Science: Photons to Phenomenology*. MIT Press, Cambridge, MA.

A tremendously informative book that deals with nearly every important empirical issue concerning the cognitive science of visual perception.

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

In contrast to Gibson's ecological approach, this book presents Rock's defense of the view that visual perception results from unconscious inference and problem solving.

Glossary terms

Sense-data -- Immediate or direct objects of acquaintance in sense perception. Traditionally believed to be private.

Disjunctivism -- Views in philosophy of perception according to which perceptual experiences and hallucinations do not belong to a common psychological kind, or share no common core content.

Representationalism -- The claim that the phenomenal characteristics of an experience ("what it's like") supervene upon its representational properties.

Transparency (diaphanousness) -- The feature of perceptual experience thanks to which attempts to introspect intrinsic qualitative features of an experience yield instead only intentional or representational characteristics of experience.

Inattentional blindness -- Failure to detect and report objects and events in a visual scene to which one does not attend, even when one might otherwise expect those objects and events to be salient.

Change blindness -- Failure to detect and report changes to a visual scene that one might otherwise expect to be salient.

Enactivism -- An approach to the nature of perception and perceptual content according to which action is constitutive of perception, and perceiving is a way of actively exploring an environment.

Blindsight -- Visual responsiveness, revealed by above chance performance, without reportable visual consciousness.