Finding out about filling-in:
A guide to perceptual completion
for visual science and the philosophy
of perception

Luiz Pessoa
Department of Computer and Systems Engineering, Center of Technology,
Federal University of Rio de Janeiro, Ilha do Fundao, Rio de Janeiro, RJ
21945-970, Brazil
pessoa@cos.ufrj.br   www.cos.ufrj.br/~pessoa

Evan Thompson[1]
Department of Philosophy and Centre for Vision Research, York University,
North York, Ontario, Canada M3J 1P3
evant@yorku.ca   www.yorku.ca/research/vision/evant.htm

Alva Noë
Department of Philosophy, University of California, Santa Cruz,
Santa Cruz, CA 95064
anoe@cats.ucsc.edu   www2.ucsc.edu/people/anoe/

Abstract: In visual science the term *filling-in* is used in different ways, which often leads to confusion. This target article presents a taxonomy of perceptual completion phenomena to organize and clarify theoretical and empirical discussion. Examples of boundary completion (illusory contours) and featural completion (color, brightness, motion, texture, and depth) are examined, and single-cell studies relevant to filling-in are reviewed and assessed. Filling-in issues must be understood in relation to theoretical issues about neural–perceptual isomorphism and linking propositions. Six main conclusions are drawn: (1) visual filling-in comprises a multitude of different perceptual completion phenomena; (2) certain forms of visual completion seem to involve spatially propagating neural activity (neural filling-in) and so, contrary to Dennett's (1991; 1992) recent discussion of filling-in, cannot be described as results of the brain's "ignoring an absence" or "jumping to a conclusion"; (3) in certain cases perceptual completion seems to have measurable effects that depend on neural signals representing a presence rather than ignoring an absence; (4) neural filling-in does not imply either "analytic isomorphism" or "Cartesian materialism," and thus the notion of the bridge locus – a particular neural stage that forms the immediate substrate of perceptual experience – is problematic and should be abandoned; (5) to reject the representational conception of vision in favor of an "enactive" or "animate" conception reduces the importance of filling-in as a theoretical category in the explanation of vision; and (6) the evaluation of perceptual content should not be determined by "subpersonal" considerations about internal processing, but rather by considerations about the task of vision at the level of the animal or person interacting with the world.

Keywords: blind spot; bridge locus; brightness; consciousness; filling-in; Gestalt theory; illusory contours; isomorphism; linking propositions; perceptual completion; vision

1. Introduction

Figure 1 illustrates the so-called neon color illusion: a red diamond is seen where there is only a lattice of red line segments. The term *neon color spreading* (van Tuijl 1975) is often used to describe this phenomenon: the figure seems to result from the color having spread between the line segments into the background. The color is said to "fill in" the background, thereby forming the figure. Nevertheless, one does not see any spreading or filling-in process; one sees only the figure. What line of reasoning, then, lies behind the description of such figures in the terminology of "filling-in"?

Visual scientists use the terms *filling-in* and *perceptual completion* to refer to situations in which subjects report

that something is present in a particular region of visual space when it is actually absent from that region, but present in the surrounding area. The idea is easiest to understand in the case of the blind spot. We have a blind spot in each eye corresponding to the region where the optic nerve leaves the retina and there are no photoreceptors (Fig. 2). In everyday perception we are never aware of the blind spot. The blind spots of the two eyes do not overlap, and stimulus that falls on the blind spot of one retina will fall outside the blind spot of the other retina. Even under monocular viewing the blind spot is not easily revealed. Close one eye and fixate on a point on a uniformly colored piece of paper: there is no experience of any gap or discontinuity in the visual field. Now follow the instructions for the blind spot demonstration in Figure 3. The left dot dis-

appears and one perceives a uniform expanse of brightness and color. This is an example of perceptual completion or visual filling-in: the color and brightness surrounding the area corresponding to the blind spot are said to "fill in" that area so that a uniform expanse is perceived.

The existence of such perceptual completion phenomena at the subject level is uncontroversial. However, the term *filling-in* is often used in a controversial sense that goes beyond what subjects report. In this controversial sense, the term *filling-in* suggests that certain kinds of subject-level perceptual completion phenomena are accomplished by the brain providing information to make up for an absence – by the brain actively filling in the missing information. However, whether there is *neural filling-in* is a matter of great debate in visual science. We will argue later that there is considerable evidence for neural filling-in, but that great care needs to be taken in thinking about the relation between neural filling-in and subject-level perceptual completion.

To appreciate the debates about filling-in it is necessary first to review certain facts about vision and second to discuss certain conceptual and methodological points. Many of the objects we perceive have roughly uniform regions of surface color and lightness. Now consider two facts. First, cells in the visual cortex in general do not respond to uniform regions, but rather to discontinuities (Hubel & Wiesel

LUIZ PESSOA is Assistant Professor of Computer and Systems Engineering at the Federal University of Rio de Janeiro, Brazil. He received a Ph.D. in Cognitive and Neural Systems from Boston University and B.Sc. and M.Sc. degrees from the Federal University of Rio de Janeiro. He has published articles on computational modeling of vision, visual psychophysics, and neural network models, among others.

EVAN THOMPSON is an Associate Professor in the Department of Philosophy and a member of the Centre for Vision Research at York University. He received his Ph.D. in philosophy from the University of Toronto. He is the author of *Colour Vision: A Study in Cognitive Science and the Philosophy of Perception* (Routledge Press, 1995), and co-author, with Francisco J. Varela and Eleanor Rosch, of *The Embodied Mind: Cognitive Science and Human Experience* (MIT Press, 1991). He has published a number of articles in philosophy and cognitive science, including a previous *Behavioral and Brain Sciences* target article on color vision.

ALVA NOË is Assistant Professor of Philosophy at the University of California, Santa Cruz. He received a Ph.D. in philosophy from Harvard University and a B.Phil. from Oxford University, and has been a Research Associate of the Center for Cognitive Studies at Tufts University. He has published articles on topics in the philosophy of perception, philosophy of mind, and other areas.

Figure 1. Neon color illusion. When the segments in the diamond are red and the matrix is black (left), the entire diamond appears reddish even though only the segments are physically colored red (right).

1962; 1968). In other words, many neurons respond more strongly to boundaries than to regions or surfaces. Second, psychophysical experiments (e.g., with stabilized images) have shown the importance of boundaries for proper surface perception (Krauskopf 1963; Yarbus 1967). In the classic study by Krauskopf (1963), an inner green disc was surrounded by a red annulus (Fig. 4). When the red-green boundary was stabilized on the retina (so that it always maintained a fixed position on the eye), subjects reported
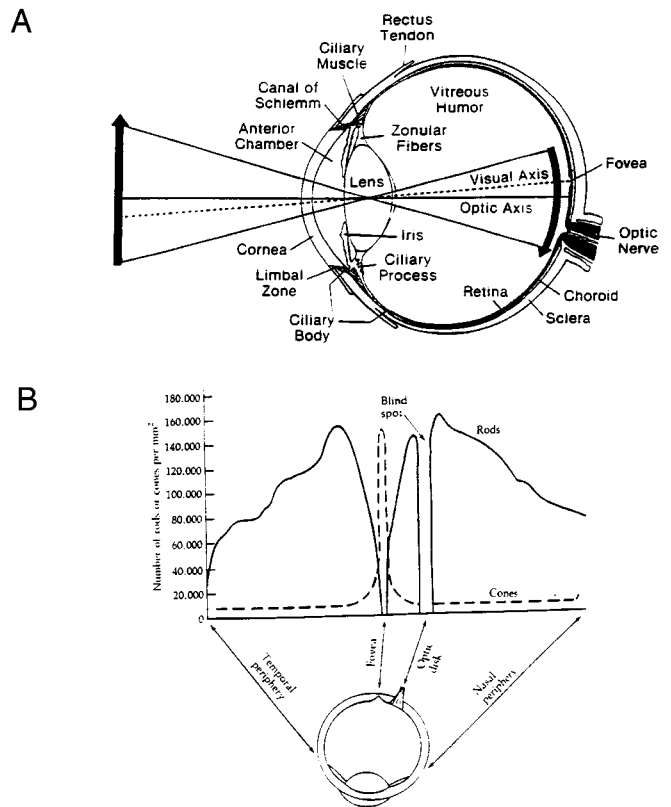

Figure 2. The blind spot. (A) Vertebrate lens and retina. The blind spot corresponds to the area where the optic nerve leaves the retina and there are no photoreceptors (Sheperd 1994; used by permission of Oxford University Press, Inc.). (B) Distribution of photoreceptors over the extent of the retina of the right eye. Note the complete absence of photoreceptors in the blind spot region (Sekuler & Blake 1985, p. 56; used by permission of the McGraw-Hill Companies).

Figure 3. Blind spot demonstration. Close your right eye while fixating at the cross. Adjust the distance of the page in front of you until the left dot disappears (start at a distance of approximately 8 inches). A similar procedure should be used with the left eye closed.

that the central disk disappeared and the whole target – disk and annulus – appeared red. (This is another case of visual filling-in: Krauskopf's observers perceived the central area as having the red of the surrounding area even though "green" light was striking the corresponding region of the retina.) These and other results (e.g., Land 1977) suggest that even under natural viewing conditions the perceived color of a surface depends not only on the light reflected from the surface but on the change in light across the boundary of the surface.

If boundaries – in fact, transients – are so important, how is the brain able to determine the color and lightness of continuous regions? Is there an active filling-in process at the neural level? Some visual scientists have developed theories and models based on the idea of a neural filling-in process that involves activity spreading, diffusion, or other forms of neural completion. Others have argued against this idea, suggesting, for example, that contrast measures at borders can be used to assign values of surface featural qualities, such as brightness or color. We will argue later that in

the case of brightness perception there is good evidence for neural filling-in that involves spatially propagating activity.

However, the filling-in controversy is not only empirical, because it involves fundamental conceptual and methodologic issues. To appreciate these issues it is necessary to introduce the concepts of the *bridge locus* and *neural–perceptual isomorphism.* The best way to introduce these concepts is by way of an example. Figure 5 presents another case often discussed in connection with filling-in, the Craik-O'Brien-Cornsweet effect (Cornsweet 1970; Craik 1940; O'Brien 1958). Two largely uniform regions of different brightness are seen, although most of the corresponding stimulus regions have exactly the same luminance. In fact, the two regions differ only in the luminance distribution at the "cusp" edge separating the two regions. Why, then, do we see a brightness step?

Here is one route leading to an answer to this question that appeals to neural filling-in (Todorović 1987). Suppose one assumes that activity of a particular type in a specific set of neurons is necessary and sufficient for the occurrence of the Craik-O'Brien-Cornsweet effect. These neurons would form the "immediate substrate" for the perceptual effect. Visual scientists use the term *bridge locus* to refer to this idea of a particular set of neurons "whose activities form the immediate substrate of visual perception" (Teller & Pugh 1983, p. 581). Now suppose one also assumes that there must be a one-to-one correspondence between the perceived spatial distribution of brightness in the effect and the

Figure 5. Craik-O'Brien-Cornsweet display. The luminance of the two rectangles is identical, except in the vicinity of the common border, in which there is a luminance cusp. For appropriate display conditions, the left and right regions appear to have uniform color and the left appears uniformly brighter than the right. Above the figure we show a one-dimensional cross section of the luminance distribution of the figure below.


Figure 4. Fading of stabilized images. (Left) The boundary between a red ring (dark gray) and a green disk (light gray) is stabilized on the retina (bottom). (Right) After stabilization a large red disk is seen (bottom). This situation should be contrasted to the case where no stabilization occurs (top).

neural activity at the bridge locus. In other words, just as the perceptual content consists of two uniform regions with a brightness step, so too the immediate neural substrate must consist of spatially continuous activity and a step difference. In short, suppose one assumes that there must be an *isomorphism* between the perceived brightness distribution and the neural activity at the bridge locus (Todorović 1987). One would then arrive at the following sort of explanation of the effect: the brain takes the local edge information and uses it to fill in the two adjacent regions so that the region with the luminance peak (left) becomes brighter than the region with the luminance trough (right). The end result is the perception of a brightness step in the absence of any corresponding luminance step.

Two basic ideas are involved here. The first is that the way things seem to the subject must be represented neurally in the subject's brain. The second is the idea that, in analyzing visual perception, one must arrive at a "final stage" in the brain – a bridge locus – where there is an isomorphism between neural activity and how things are perceived by the subject. (The isomorphism can arise at some earlier stage of visual processing, as long as it is preserved up to the bridge locus; see Teller 1984, p. 1242; Teller & Pugh 1983, p. 586; Todorović 1987, p. 550.) We will refer to this idea as *analytic isomorphism.* When applied to perceptual completion phenomena, such as the Craik-O'Brien-Cornsweet effect, analytic isomorphism entails that there *must* be neural filling-in to make up the difference between how things are and how they are perceived by the subject.

Analytic isomorphism is essentially a conceptual or methodologic doctrine about the proper form of explanation in cognitive neuroscience. The doctrine is that it is a condition on the adequacy of an explanation that there be a bridge locus where an isomorphism obtains between neural activity and the subject's experience. Furthermore, the isomorphism typically is assumed to hold for spatial or topographic properties, thus suggesting that vision involves representations having the form of an "internal screen" or "scale model" that preserves the metric properties of the external world (O'Regan 1992). In this target article, we argue that analytic isomorphism should be rejected. Nevertheless, we believe that the *empirical* case for neural filling-in remains strong.

Enter the philosophers. Dennett (1991; 1992) has tried to brand "filling-in" the "F-word" of cognitive science. He thinks that the sort of reasoning epitomized by analytic isomorphism, and hence the idea that there must be neural filling-in, depend on a fundamentally mistaken conception of consciousness. Dennett calls this conception *Cartesian materialism.* In the stereotypical version of Cartesian materialism, there is a place in the brain – a "Cartesian theater" – in which contents become conscious as a result of being presented to an inner "audience" or homunculus – a viewer of the panoramic "internal screen" (O'Regan 1992). Most scientists agree that this idea is totally wrong. Nevertheless, Dennett thinks that not everyone understands exactly why it is wrong. In Dennett's assessment, the real mistake is a *conceptual* one: the mistake is to assume that consciousness is a property of individual contents in the way that truth can be considered a property of individual sentences. Given this concept of consciousness, it would seem that there must be a determinate spatio-temporal point in the brain where a content "enters consciousness." Dennett

thinks that this concept of consciousness is incoherent and offers in its place the idea that "consciousness is a species of mental fame" (Dennett 1996b). Just as it is impossible to be famous for a second, or to become famous in a second, or to be famous when there are no other people around, so it is impossible for a single, momentary, isolated content to become conscious: for a content to become conscious it has to persist long enough to achieve certain effects on memory and the control of behavior. In short, for Dennett, consciousness is constituted through the joint interaction of spatially and temporally distributed information-processing systems.

Dennett argues that to claim there is neural filling-in "is a dead giveaway of vestigial Cartesian materialism" (1991, p. 344). Like a number of visual scientists, he believes there is no reason to suppose that the brain fills in the regions; the brain simply represents the fact that regions are filled-in without itself doing any filling in. For Dennett, perceptual completion is a case of the brain "finding out" or "judging" that certain features are present, without the brain having to "present" or fill in those features.

In visual science, O'Regan (1992) has outlined a similar position. He suggests that the need to appeal to neural filling-in would "evaporate if we abandon the idea that 'seeing' involves passively contemplating an internal representation of the world that has metric properties like a photograph or scale model" (1992, p. 483). Instead, he argues that "seeing constitutes an active process of probing the external environment as though it were a continuously available *external memory*" (p. 484). Seeing depends not on the filling-in of a metric representation, but rather on "interrogating" the external environment directly through eye movements, and then integrating the altered retinal sensations into one's cognitive framework (p. 475).

In the past few years, the amount of literature on filling-in among psychologists, neuroscientists, theoretical modelers, and philosophers has increased. However, the term *filling-in* continues to be used in different ways. Sometimes it is used to describe what the subject perceives; sometimes it is used to refer to what the brain does. The term is also used to describe different sorts of perceptual completion. For example, although illusory contours and brightness perception probably involve different processes, the concept of filling-in is often used in association with both: a line segment is said to fill in between the inducers, and brightness is said to fill in across regions. Although one need not argue which usage is preferable, it should be obvious that without conceptual and terminological clarification there is room for considerable confusion. For example, do these two types of completion involve common principles and mechanisms? Or are they distinct? Given this situation, it can be hardly surprising that in the more theoretical debates the participants often seem to be talking past one another.[2]

In this target article we provide a taxonomy of perceptual completion phenomena with an overview of some recent psychophysical results pertinent to filling-in. The taxonomy is meant as a step toward conceptual and terminological clarification. We would like to emphasize at the outset that our taxonomy is based on salient examples of perceptual completion; it is in no way an exhaustive survey.

In addition we argue that the filling-in issues are best understood in relation to issues about neural–perceptual isomorphism and "linking propositions." Linking propositions

are statements that relate perceptual states to physiological states (Teller 1980; 1984; 1990; Teller & Pugh 1983). The concept of a neural–perceptual isomorphism relies on a certain sort of linking proposition, to be discussed later. Nonisomorphic approaches reject the idea that there has to be any structural one-to-one correspondence between perceptual experience and neural processes.

Dennett's discussion of Cartesian materialism is directly relevant to these issues. Dennett has done a service by showing how the filling-in idea often depends on Cartesian materialism (see also O'Regan 1992), but we disagree with his positive view that perceptual completion is always just a matter of the brain's "finding out." We think that the idea of neural filling-in has to be separated from Cartesian materialism and analytic isomorphism. We will show that there is evidence in visual science to support the idea of neural filling-in. As discussed below, filling-in is not always just finding out. Nevertheless, we agree that it is a mistake to invoke filling-in as a theoretical category in the explanation of vision. As we see it, such invocations depend on a mistaken representational account of the task of vision. We argue later that the representational account should be replaced by the sort of account variously described as "animate" (Ballard 1991; 1996; Ballard et al. 1997) or "enactive" (Thompson 1995; Thompson et al. 1992; Varela et al. 1991).

The next section discusses in more detail the notion of linking propositions and the concept of neural–perceptual isomorphism. Section 3 offers a taxonomy of filling-in phenomena to help clarify and organize the empirical findings, plus a detailed look at examples of boundary completion involving illusory contours, and featural completion involving color, brightness, motion, texture, and depth. Section 4 examines some neurophysiological data about perceptual completion. Section 5 summarizes Dennett's position on filling-in and connects it to earlier discussions of isomorphism in visual science. Section 6 presents evidence for neural completion that is inconsistent with Dennett's position. Section 7 reviews some studies that assess the measurable effects of perceptual completion. Section 8 shows how neural filling-in does not entail either Cartesian materialism or analytic isomorphism. Section 9 introduces the personal/subpersonal distinction – the distinction between, on the one hand, the perceiving animal or person as a whole interacting with its environment, and on the other hand, the animal's internal functional organization and processing – and shows its relevance to the filling-in controversy. Section 10 concludes with a statement of directions for further research.

2. Linking propositions and isomorphism

In this section we briefly review some important conceptual and methodologic issues about explanation in visual science.

In 1865, Mach stated what has since become known as "Mach's principle of equivalence":

> Every psychical event corresponds to a physical event and vice versa. *Equal* psychical processes correspond to *equal* physical processes, *unequal* to *unequal* ones. When a psychical process is analyzed in a purely psychological way, into a number of qualities *a, b, c,* then there corresponds to them just as great a number of physical processes α, β, γ. To all the details of psychological events correspond details of the physical events (Mach 1865/1965, pp. 269–70).

Thirteen years later in 1878 Hering (1878/1964) asserted that the neural–perceptual parallelism was a necessary condition of all psychophysical research. Müller (1896) then gave a more explicit description of the neural – perceptual mapping. He proposed five "psychophysical axioms" that postulated a one-to-one correspondence between neural and perceptual states (see Scheerer 1994; Teller 1984). In particular, his second axiom stated that perceptual equalities, similarities, and differences correspond to neural equalities, similarities, and differences. This axiom was not offered as a solution to the so-called mind-body problem, but rather as a methodologic principle that could be a guide in inferring neural processes from perceptual experiences (Scheerer 1994, p. 185).

Köhler accepted this idea, but thought that Müller's axioms were not comprehensive enough because they did not include occurrent perceptual states, but covered only the logical order between neural and perceptual states (Scheerer 1994, p. 185). In 1920 he proposed what he would later call the principle of isomorphism (Köhler 1920), building on Müller's earlier formulation, as well as Wertheimer's (1912). In his 1947 book *Gestalt Psychology,* Köhler wrote: "The principle of isomorphism demands that in a given case the organization of experience and the underlying physiological facts have the same structure" (Köhler 1947, p. 301).

There are several points about Köhler's principle of isomorphism that deserve mention. First, by the phrase "have the same structure" Köhler had in mind structural properties that are *topological.* Although the concept of neural–perceptual isomorphism has often been taken to mean a geometric one-to-one mapping, Köhler clearly intended the isomorphism concept to have a topological sense. For example, he argued that spatial relationships in the visual field cannot correspond to geometric relationships in the brain; they must correspond rather to functional relationships among brain processes (Köhler 1929, pp. 136–41; 1930, pp. 240–49).

Second, Köhler did not hypothesize that neural–perceptual isomorphism was valid for all properties of perceptual experience. In particular, he did not extend the principle of isomorphism to sensory qualities, such as brightness and color (Köhler 1969, pp. 64–66, as quoted in Kubovy & Pomerantz 1981, p. 428). The principle was restricted to "structural properties" of the perceptual field, that is, to characteristics of perceptual organization, such as grouping and part-whole relationships.

Finally, it is not clear whether Köhler espoused what we are calling analytic isomorphism. Two considerations suggest that he did not. First, Köhler upheld a nonlocalizationist view of brain function, in which field physics was the main analogy for the underlying physiology; hence the notion of a privileged site of perceptual experience in the brain seems foreign to his way of thinking about the neural–perceptual relation.[3] Second – and this is more telling – he seems to have held (at least according to one interpreter) that the isomorphism principle "is not an a priori postulate, but 'remains an hypothesis which has to undergo one empirical test after the other' " (Scheerer 1994, p. 188).

In 1969, Weisstein provided a clear discussion of the relationship between neural states and perceptual states in the context of experimental studies based on recordings from single cells: "Axiomatically, it can be assumed that any visual event has some corresponding neural circuitry, and

that a good deal of neural circuitry in the visual system has a corresponding function in producing a perceptual event" (Weisstein 1969, p. 159). She noted that it is important "to choose some aspect of the behavior shown in single units which appears to have an analogous psychophysical effect in humans," but also that "the corresponding psychophysical effect cannot be strictly analogous to a single unit recording," for there will be more than one single unit activated for almost any conceivable stimulus. As Weisstein rightly insisted, both a spatial and a temporal characterization of neural and perceptual data are essential in establishing a bridge between the two domains that goes beyond identifying rough similarities (see sect. 4.5).

In recent years Teller has reintroduced some of these issues into visual science (Teller 1980; 1984; 1990; Teller & Pugh 1983). According to Teller, acceptable explanations within visual science have the following form:

> If the question is, what is it about the neural substrate of vision that makes us see as we do, the only acceptable kind of answer is, we see X because elements of the substrate Y have the property Z or are in the state S (Teller 1990, p. 12).

This formulation leaves open another question about form: what is the relation between the form of a given neural response and the form of the corresponding visual appearance? Answers to this question invoke linking propositions – propositions that relate neural states to perceptual states. By analyzing how visual scientists reason, Teller (1984) formulated five families of linking propositions called *identity, similarity, mutual exclusivity, simplicity,* and *analogy.*

The analogy family is the one that concerns us here. It is a "less organized" family of propositions whose form is as follows:

$$\phi \text{ "looks like" } \psi \rightarrow \phi \text{ explains } \psi,$$

where ϕ stands for physiologic terms, and ψ stands for perceptual terms. The arrow connective $\rightarrow$ has a conditional sense, thus the formulation reads: If the physiologic processes (events, states) "look like" the perceptual processes (events, states), then the physiologic processes explain the perceptual processes.

The arrow is not the connective of logical entailment. It is heuristic and is meant to guide the search for the major causal factors involved in a given perceptual phenomenon. Thus the term "explains" on the right-hand side is really too strong – the idea is that ϕ is the major causal factor in the production of ψ: "if psychophysical and physiological data can be manipulated in such a way that they can be plotted on meaningfully similar axes, such that the two graphs have similar shapes, then that physiological phenomenon is a major causal factor in producing the psychophysical phenomenon" (Teller 1984, p. 1240).

The analogy family of linking propositions is similar to Köhler's principle of isomorphism but more general. Isomorphism in Köhler's sense can be seen as a particular instance of the analogy idea, one in which "looks like" is taken in the sense of structural correspondence. In visual science today, this idea of a one-to-one structural correspondence is often taken to mean a spatial correspondence, so that, for example, spatial variations of brightness in the visual field are explained by analogous spatial variations of neural activity (Todorović 1987, p. 548).

To argue that the brain does not really fill in, it only "finds out," means that one rejects the principle of isomorphism applied to perceptual completion. In other words, one rejects the hypothesis that perceptual completion depends on neural completion processes that are structurally isomorphic to the perceptual phenomena. We agree that the doctrine of analytic isomorphism should be rejected, but we think there is evidence for neural filling-in. One problem with Dennett's treatment is that he applies his "filling-in is finding out" point across the board, without considering the different kinds of perceptual completion. Before we go further, then, we need to review the different sorts of perceptual completion phenomena.

3. Filling-in: A taxonomy of perceptual completion phenomena

The following working taxonomy is meant as a step toward conceptual and terminological clarification. The term *perceptual completion* refers to what subjects report and should be taken in a theory-neutral sense. It is not meant to have any implications for whether there are, or are not, neural filling-in mechanisms in the operation of the visual system. This is a matter to be taken up later. All that the term *perceptual completion* is meant to imply is that subjects report that something seems to be present in a particular region of visual space when it is actually absent from that region, but present in the surrounding area.

There are two general divisions in the classification: (1) amodal completion versus modal completion and (2) boundary completion versus featural completion. These two divisions cross-classify each other, and, thus, there is no hierarchical organization implied in this listing. We will introduce each division briefly using examples and then consider them in depth in the discussion section.

3.1. Amodal completion and modal completion

Michotte et al. (1964) distinguished between two types of perceptual completion, *modal* and *amodal.* In modal completion, the completed parts display the same type of attributes or "modes" (e.g., brightness) as the rest of the figure. Illusory figures provide a particularly interesting example. Figure 6 shows the famous Kanizsa triangle (Kanizsa 1955; 1979). Here there are illusory contours – clear boundaries where there is no corresponding luminance gradient – and a brightening within the figure. The illusory contours and the central brightening are modal in character: they are perceptually salient and appear to belong to the figure rather than the ground.[4]

Amodal completion refers to the completion of an object that is not entirely visible because it is covered or occluded by something else (Kanizsa & Gerbino 1982). Thus, amodal completion denotes the perception of parts of objects – the completed regions – that entirely lack visible attributes. For example, consider Figure 7: although the circles are occluded, they are easily recognized and are seen as lying underneath the rectangles. The parts of the circles occluded by the rectangles are said to be *amodally present.*

3.2. Boundary completion and featural completion

The distinction between boundary completion and featural completion was first proposed in the theoretical work of Grossberg and Mingolla (Grossberg 1987a; 1987b; Grossberg & Mingolla 1985).[5] Illusory figures also can serve to

Figure 6. The Kanizsa triangle. Illusory contours are seen forming a triangle-shaped region although there are no corresponding luminance changes. Note also that the illusory figure is brighter than the background.

introduce the distinction. In the Kanizsa triangle, there is boundary completion – the illusory contours complete to form a triangular outline – and there is featural completion, an illusory brightening within the figure compared with the background in the absence of any luminance difference.

In the next section (sect. 3.3) we review both boundary and featural completion phenomena. Before doing so we wish to state explicitly that we are not at all committed to the currently popular "feature-based" approach to visual perception, that is, to the idea that features such as color, brightness, texture, and so on are the "visual primitives" out of which visual perception is composed. Our intention is not to subscribe to this paradigm; it is simply to review some psychophysical studies that provide evidence for various sorts of perceptual completion, and these studies typically focus on the visual attributes just mentioned.

3.3. Discussion

3.3.1. Amodal versus modal completion. Amodal completion is the perceptual completion of occluded objects; modal completion is perceptual completion in the foreground. What does it mean to say that the amodally present parts of the figure are "seen" or "recognized"?

Kanizsa and Gerbino (1982) have explored these questions in terms of the relation between seeing and thinking. They describe the amodal presence of the occluded parts as having an "encountered" character (their translation of Metzger's *angetroffen*), and they say that "the name 'amodal presence' is reserved for the 'encountered' presence of parts not directly visible" (pp. 171–72). They contrast this encountered presence with a "purely mental completion of an inferential kind," saying that "amodal completion transforms a collection of pieces into a reality of complete things of a phenomenal 'encountered' character" (p. 173). Kanizsa and Gerbino also try to differentiate between "cognitive completion" and "perceptual completion" by arguing that the latter always has a functional effect on the visual aspect of a situation, whereas the former does not suffice to produce such effects. Of course, these two types of completion are not mutually exclusive; they can act together. However,


Figure 7. Amodal completion. (Top) Four disks are easily recognized as being occluded by four rectangles. This percept should be compared with the one on the bottom from which the occluding rectangles have been removed. Adapted from Kanizsa and Gerbino (1982).

it seems they can be dissociated, too: the cognitive sort can be present in the total absence of the perceptual sort (pp. 173–74).

Cognitive explanations have also been proposed for modal completion phenomena. For example, during the 1970s, the preferred explanation of illusory contours involved appealing mainly to cognitive-like processes of postulation and hypothesis formation (Gregory 1972; Rock & Anson 1979) (see sects. 4.1, 6.1, 7.3).

We think that a profitable approach to these issues would

be to determine to what extent modal and amodal completion involve common mechanisms. For example, in a series of investigations with illusory contours and occluded figures (among other stimuli), Kellman and Shipley (1991) have gathered evidence that common interpolation mechanisms are involved. On the other hand, Anderson (1995) suggests that two kinds of boundary interpolation are involved in modal and amodal completion phenomena.

Another way to probe completion mechanisms is to test current theories of surface perception. For example, Grossberg's FACADE theory tries to explain many challenging completion phenomena involving form, color, and depth (Grossberg 1994).

3.3.2. Amodal completion, surface interpolation, and the blind spot. Several studies of the blind spot have appeared recently (Brown & Thurmond 1993; Komatsu & Murakami 1994; Murakami 1995; Tripathy et al. 1995; Tripathy & Levi 1994). Here we mention a study related to amodal and modal perceptual completion. Durgin et al. (1995) argue that filling-in of the blind spot be considered a case of surface interpolation. They compared percepts involving the blind spot to percepts involving occlusion (e.g., a disk lying on a thick line). In all the tasks they investigated, including motion stimuli and amodal completion, the percepts were similar in detail.

Durgin et al. interpret their results as "consistent with the null hypothesis that the blind spot is treated visually as a region of little or no information" (1995, p. 837). Although the statement in this particular form is uncontroversial – there are no photoreceptors originating from the blind spot region – Durgin et al. go on to state: "we do not consider that the content of visual perception in the blind spot must directly reflect the activity of 'filled in' visual maps" (p. 837).

Durgin et al.'s perceptual demonstrations support the notion that the blind spot may be treated as an "occluded" region of vision "without an occluder." However, at present their demonstrations do not speak to the nature of the underlying neural processes. Durgin et al. propose that blind spot completion is similar to amodal completion, but the perceptual similarities between blind spot completion and amodal completion are not enough to draw conclusions about the neural processes involved in blind spot filling-in.

3.3.3. Boundary completion versus featural completion. Experimental studies in psychophysics provide evidence for the existence of two separate types of perceptual completion: boundary completion and featural completion. Two examples can be given, illusory contours and neon color spreading.

3.3.3.1. Illusory contours. There are three defining properties of illusory contours: clarity (or sharpness of the contours), brightness (of the illusory figure), and depth (the "depthfulness" of the illusory figure) (Lesher 1995). Illusory figures need not exhibit all of these properties. For example, illusory contours can appear without an accompanying illusory figure (no depth), such as in offset-grating stimuli. Most important in the present context are illusory figures without clarity (Fig. 8a) and illusory figures with clarity but without any accompanying brightening or darkening effect (Fig. 8b). (Another example of clarity without brightness is the phenomenon of spontaneously splitting figures: Kellman & Shipley 1991; Koffka 1935; Petter

Figure 8. Illusory contours. (A) Illusory figure without clarity (no clear contours) but with noticeable brightening. (B) Illusory figure with clarity but no brightness difference (actual appearance may be modified by reproduction processes).

1956.) We find these two situations (Fig. 8) useful for illustrating the distinction between boundary completion and featural completion. In particular, the relative independence of contour clarity and figure brightness (see Lesher 1995) suggests that independent neural mechanisms subserve boundary completion and featural completion (Grossberg & Mingolla 1985).

3.3.3.2. Neon color spreading. Watanabe and Sato (1989) studied neon color spreading in the Ehrenstein-plus-cross configuration. They varied the luminances of the cross and the outer segments (Ehrenstein inducers) and were able to show conditions in which no illusory contours formed while

color spreading occurred. Illusory contours did not occur when the cross (yellow) and the outer segments (white) were equiluminant. In these conditions color spreading was possible. Watanabe and Sato conclude that separate mechanisms subserve the two aspects of the neon color spreading effect. In their view, a luminance difference between the outer segments and their surroundings (the inner segments plus the background) induces the illusory contours. Color spreading depends on the color difference between the inner and outer segments.[6]

3.3.4. Boundary completion. Two major types of illusory contours can be distinguished: edge-induced and line-induced. Edge-induced illusory contours consist of solid inducing elements containing edges, or gaps, locally consistent with an occluding figure of the same luminance as the background. In such cases, the illusory contour is colinear with the inducer edges consistent with the occlusion. Line-induced illusory contours, on the other hand, can be seen as the limiting case of edge-induced figures, where the inducers are typically "thin." In this case, the associated illusory contours are not parallel to the inducers, but instead roughly perpendicular to them. Figure 6 (the Kanizsa triangle) presents both types of inducers. The three black, circular "pac-men" act as edge inducers, while the thin lines work as line-end inducers. Figure 8 presents two other line-end–induced illusory figures. Figure 8b also illustrates that there is no sharp separation between edge– and line-end–induced illusory contours – a line always has some width and hence provides some edge information for very small spatial scales.

The determinants of illusory contour strength are varied and include both low- and high-level factors. Low-level factors include the spatial extent and proximity of inducers, number of inducers, inducer luminance and contrast, and inducer alignment. High-level factors include perceptual set and memory, depth modifications, and inducer completeness (Lesher 1995). The strength of edge– and line-end–induced illusory contours is enhanced by providing additional stereo and motion cues. Remarkably, illusory contours can be generated *without* any luminance discontinuities whatsoever.

Random dot stereograms (Julesz 1971) give rise to strong, sharp illusory contours at the (depth-induced) edges. In such cases, the contours are associated with local depth disparity cues. Illusory contours also occur without local disparity (and luminance) discontinuities. If only the edge inducers of a Kanizsa figure are defined by disparity information in a random dot stimuli, contour completion ensues (Mustillo & Fox 1986; Prazdny 1985). In this case, stereo cues give rise to inducers, which then behave as luminance defined inducers, resulting in contour formation across regions with no local stereo (and luminance) cues. Analogously, illusory contours can be specified by local motion cues. (See Lesher 1995 for an in-depth treatment of the above issues.)

Illusory contour boundary formation is accompanied by modal completion – the presence of perceptually salient figure-like (as opposed to ground-like) elements. Amodal completion also involves boundary completion. In other words, contours are formed in regions that are not visible. Although many studies of amodal completion involve both the completion of contours and the completion of regions, we include amodal completion in this section because

several of the issues related to boundary formation (e.g., good continuation) are critical for the study of amodal completion.

Amodal completion is governed by global and local factors. Global factors include symmetry and simplicity (Buffart et al. 1981; 1983). Defenders of the primacy of global factors see the whole figure as important in determining the percept, which is often postulated to be the simplest organization of the stimulus. Local factors include contour continuity and curvature (Shipley & Kellman 1992). Those who espouse the importance of local factors stress the autonomy of relatively low-level contour processing mechanisms as governing amodal completion. The debate continues on the relative merits of global (Buffart et al. 1981) and local (Boselie 1994; Kellman & Shipley 1991; Takeichi et al. 1995) factors in determining amodal completion, with several hybrid schemes having been proposed (Boselie 1988; Sekuler 1994; Wouterlad & Boselie 1992). Recent studies have also attempted to dissect potential representational stages involved in amodal completion, such as a mosaic stage – in which a literal description of the visible parts of an occluded surface (called *mosaic*) would be produced – and a completion stage (Bruno et al. 1996; Sekuler & Palmer 1992). Finally, as previously discussed, current research targets the critical issue of whether modal and amodal completion phenomena are subserved by common mechanisms.

3.3.5. Featural completion. For what kinds of feature can there be modal perceptual completion? Experimental studies indicate that featural completion can involve brightness, color, texture, motion, and depth.

3.3.5.1. Brightness and color. We have already seen examples of completion for brightness and color – illusory figures and neon color spreading, as well as the Craik-O'Brien-Cornsweet effect discussed in section 1. In general, brightness and color filling-in are related to the role played by contours in surface perception. The nature of the underlying neural mechanisms is a matter of considerable debate. Several researchers have advanced models based on the idea of filling-in at the neural level (Arrington 1996; Cohen & Grossberg 1984; Davidson & Whiteside 1971; Fry 1948; Gerrits et al. 1966; Gerrits & Vendrik 1970; Grossberg & Todorović 1988; Hamada 1984; Neumann 1996; Pessoa et al. 1995; Ross & Pessoa 1997; Walls 1954). Other researchers have proposed mechanisms that do not actively fill in regions, but use other processes to assign feature (see the discussion in Kingdom & Moulden 1989; Pessoa 1996a; Ratliff & Sirovich 1978).

3.3.5.2. Texture. Watanabe and Cavanagh (1991; 1993) investigated whether attributes such as texture can make up an illusory surface in a configuration similar to one that elicits neon color spreading. Subjects viewed a display consisting of crosses filled with black and white textures inserted into the gaps of the Ehrenstein figure (Fig. 9). They were asked to report "whether the texture of the crosses appeared to spread outside of the cross" or not; 83.3% of the subjects reported the perception of texture outside of the cross, compared with 0% when the central cross regions were viewed in isolation. In the case of neon color spreading, the spreading effect decreases when the inner cross is disconnected from the Ehrenstein figure – for example,

Figure 9. Texture filling-in displays investigated by Watanabe and Cavanagh (1991). (A) Crosses made of black and white diagonal stripes inserted into the gaps of an Ehrenstein figure. When foveating the center of the grid, subjects perceive a circular contour around the texture crosses. A large percentage of the subjects also report texture within the circular disks. (B) The texture crosses presented alone. Neither circular contours nor texture spreading is reported in this case (Watanabe & Cavanagh 1993; used with permission).

when the cross is rotated around its center (Redies & Spillmann 1981). Watanabe and Cavanagh report that their demonstrations with texture spreading show the same tendency. However, unlike neon color spreading, texture spreading totally disappeared when a textured cross was foveated (which the reader can confirm in Fig. 9a).

Studies by Kawabata (1982; 1984; 1990) also provide evidence about the perceptual completion of texture. Kawabata (1982) investigated the stimulus conditions necessary for perceptual completion across the blind spot. Several types of pattern were investigated, including dotted lines (each half to one side of the blind spot), parallel lines (grating pattern), and concentric circles. Grating patterns completed only when they covered two quadrants around the blind spot; for example, covering the top hemifield or the left hemifield. Concentric circle patterns also completed across the blind spot as long as the pattern consisted of more than three circles. Perceptual completion of dotted lines occurred as long as there was a small spatial separation between the line ends and the borders of the blind spot; otherwise they were perceived as two independent lines. In a subsequent study, Kawabata (1990) showed that types of completion similar to those obtained for the blind spot were obtained in peripheral vision (15 degrees or more).

A recent series of studies by Ramachandran and colleagues on perceptual completion effects has also attracted considerable attention (Ramachandran 1992a; 1992b; 1993a; Ramachandran et al. 1993; Ramachandran & Gregory 1991). In one study, a homogeneous gray (or pink) square was displayed on a dynamic two-dimensional noise pattern (twinkling noise). Three effects were reported: (1) the gray square faded after approximately 5 seconds; (2) the square region appeared filled in with the surrounding noise pattern; and (3) the noise pattern in the region originally occupied by the gray square persisted. Ramachandran and colleagues described (2) as a filling-in effect and discussed possible neural substrates and mechanisms (see Ramachandran et al. 1993). Finally, Hardage and Tyler (1995) recently compared the actual filling in (Ramachandran's

second effect) with the twinkle aftereffect (the third effect). They showed that the two effects are sensitive to different spatial and temporal parameters. Their results indicate that different mechanisms are involved in the generation of these two percepts.

3.3.5.3. Motion. Watanabe and Cavanagh (1993) (in the study previously mentioned) also employed the Ehrenstein figure-paradigm to investigate motion spreading. The texture inside the crosses was made to move while the crosses themselves remained stationary. All of their subjects reported that motion appeared to spread outside the crosses.

Apparent motion (the phi phenomenon) is a well known effect. Motion perception is induced by stimuli presented at distinct spatial positions and with an appropriate temporal interval between them. Lockhead et al. (1980) investigated a variation of apparent motion known as *sensory saltation* (Geldard & Sherrick 1972). In this perceptual effect (for proper display parameters) the subject reports discrete points of stimulation between the two actual stimulation points. Lockhead et al. attempted to determine whether illusory stimulation points in sensory saltation could be spatially assigned to the (receptorless) blind spot region. The results for three subjects were similar and showed that the saltation crossed the blind spot, with the illusory stimulation points often localized within the blind spot region.

3.3.5.4. Depth. The completion of depth is vividly illustrated by the stereogram shown in Figure 10 (Nakayama & Shimojo 1990a). For such an untextured figure, it is not clear what classical stereopsis would predict for the perception of the center of the cross. Should the percept assume the depth of the nearby vertical segment, or that specified by the horizontal limbs? For the fusion of left and center images, although the vertical limb is closer to the center, one perceives a horizontal bar in front of a vertical bar, illustrating that disparity information can "propagate" when necessary.

Another example of depth completion is given by the phenomenon of Da Vinci stereopsis (Gillam & Borsting

Figure 10. Cross stereogram. Fusion of the left and center images (by "uncrossing") produces the percept of a black horizontal bar in front of a vertical bar (as indicated on the bottom left). Fusion of the center and right images produces the percept of a vertical bar in front of a horizontal bar (bottom right). For crossfusers the reverse depth relationships are perceived. Adapted from Nakayama and Shimojo (1990).

A BC D A B C D

Left eye view   Right eye view

Figure 11.   Da Vinci stereopsis. Example of Da Vinci stereogram derived from viewing a three-dimensional scene of a room. Region BC is seen monocularly (in the right eye view only). This diagram represents the retinal images impinging on an observer near an occlusion in depth caused by a nearer object (surface AB). Adapted from Grossberg (1994).

1987; Nakayama & Shimojo 1990b). When we view a farther surface that is partly occluded by a nearer surface, one eye typically registers more of the farther surface than the other eye does. The perception of the farther surface is often derived from the view of the eye that registers more of this surface. Da Vinci stereopsis is illustrated in Figure 11. Observers see the right-eye view of the surface BD in depth, although the region that lies between the vertical lines B and C is registered monocularly by only the right eye. An outstanding question in the study of three-dimensional vision is how the monocularly viewed region BC inherits the depth of the binocularly viewed region CD (Grossberg 1994; Grossberg & McLoughlin 1997).

A final example is a stereogram by Julesz (1971, p. 336) in which each image contains 5% black dots on a 95% white background. A portion of the black dots has disparity, while the remaining ones have zero disparity. When the left and right images are stereoscopically viewed, the black dots with disparity are, as expected, seen in front. These black dots, however, cause the white surround that they enclose to be seen, as a whole, as a planar surface lying in front of another planar surface containing the zero disparity black dots and the white region that they enclose.

3.3.6. Summary. We offer a taxonomy of perceptual completion whose divisions are amodal versus modal completion, and boundary versus featural completion. Within these divisions there are numerous different sorts of perceptual completion and care must be taken to distinguish among them. In particular, there seems to be an important difference between boundary completion and featural completion, with featural completion occurring for brightness and color, texture, motion, and depth. This means that propositions about filling-in as a whole are of limited use and should be greeted with suspicion. For example, in advance of further research, there is no reason to group together illusory contour formation, for which there is strong evidence for neural completion (see sects. 4.1, 6.1, and 7.3.2), and the texture completion studies by Watanabe and Cavanagh (1991) and Kawabata (1990), in which peripheral

stimulation is necessary to elicit the completion effects. Indeed, without additional careful experiments to assess the *effects* of perceptual completion (see sect. 7), it is not clear what these studies reveal about perceptual completion at the neural level.

4. Cortical data and perceptual completion

We now turn to examine some recent physiological data relevant to perceptual completion. All the data come from single-cell studies of the response properties of cortical neurons. We wish to state here that our aim is to examine some important representative studies, not to give an exhaustive review.

4.1. Illusory contours and cortical cell responses

In an influential paper, von der Heydt et al. (1984) presented results from single-cell recordings suggesting neural correlates of illusory contours in area V2 of the macaque monkey. Almost half the cells examined exhibited sizeable responses to drifting bars or edges and also to the illusory contour induced by drifting line gratings. Cells were not simply responding to individual line-ends, however, because the typical cell would not respond to a grating with only two or three bars, but would respond with increasing strength as other bars were added, until a saturated level of activity was reached.

Von der Heydt et al. (1984) also studied neural responses to notch stimuli – dark rectangles with parts missing that formed an illusory rectangle. Cellular activity fell off with increasing notch separation and was greatly reduced when only a single notch was present, in parallel with the perceptual disappearance of the illusory figure. In all, the cellular recordings of von der Heydt et al. revealed cells whose responses to illusory contour variations resembled human psychophysical responses to similar variations (see also Redies et al. 1986 for similar results in cat visual cortex). Although some have described these findings as the dis-

covery of "illusory contour cells" (Lesher 1995), von der Heydt et al. (1984) tried to draw a clear distinction between the stimulus-response relationship, on the one hand, and perceived entities, on the other. For instance, they used the term *illusory contour stimuli*, rather than *illusory contour cells*, and they borrowed the term *anomalous contours* from Kanizsa (1955; 1979) to define a stimulus property without reference to perception.

In a recent study, Grosof et al. (1993) suggested the existence of neurons in V1 of macaque monkeys that respond to line-end stimuli similar to offset gratings. However, the results remain controversial and await further controls to establish the role these neurons play in illusory contour perception (see Lesher 1995).

4.2. The blind spot and receptive field dynamics

Fiorani et al. (1992) recorded the activity from single neurons in parts of area V1 that correspond to the blind spot. A cortical region that corresponds to the blind spot for the contralateral eye (the eye on the opposite hemispheric side) will also correspond to a normal area of the retina for the ipsilateral eye (the eye on the same hemispheric side). How do neurons corresponding to the blind spot for the contralateral eye respond when a stimulus is presented across this area and the ipsilateral eye is closed? One might predict that the neurons would not respond at all, that they would respond only to stimuli presented to the nonblind region of the ipsilateral eye. However, this is not what happens to 20% of the cells. In other words, some neurons do respond to stimuli presented through the blind spot. Especially interesting is a subpopulation of cells that Fiorani et al. call *completion neurons.* These neurons, whose receptive fields are located inside the blind spot, respond to a bar longer than the diameter of the blind spot when it is swept across the blind spot, but they respond poorly or not at all to bars restricted to one side of the blind spot. Because these neurons retinotopically map the blind spot (of the contralateral eye) where there are no receptors, Fiorani et al. describe them as having "interpolated receptive fields" (see Fig. 12).

4.3. Scotomata and receptive field dynamics

Two studies involving dynamic changes in receptive field size in primary visual cortex come from Gilbert and colleagues (Gilbert & Wiesel 1992; Pettet & Gilbert 1992; see also Gilbert 1992). Gilbert and Wiesel (1992) recorded activity from neurons in V1 in the monkey both before and after retinal lesions. They found that, over a period of minutes, neurons whose receptive field centers are originally located near the edge of the retinal scotoma have greatly enlarged receptive field sizes. They also found that, over a period of 2 months, neurons with receptive fields originally located within the lesioned area regain visual activity, but now corresponding to retinotopic positions outside of the lesioned area, thus providing an enlarged representation of the area surrounding the retinal scotoma. In other words, the retinal lesion temporarily silences cortical areas – in effect creating a cortical scotoma – but over a period of months the cortical area regains activity due to the dynamic changes in the receptive fields.

In a second study Pettet and Gilbert (1992) masked areas covering receptive fields of neurons in V1 of the cat,


GOOD RESPONSES

POOR RESPONSES

Figure 12. Completion neurons and interpolated receptive fields. Cell responses in the optic disk representation to oriented bars (left, vertical; right, horizontal). Stars and dashed lines indicate locations where stimulation with long bars elicited responses. Arrows indicate direction of motion of oriented black bars. IRF: interpolated receptive field (black squares). Adapted from Fiorani et al. (1992).

thereby creating an artificial scotoma. They found that when the area of the visual field surrounding the scotoma is stimulated, the receptive field can expand an average of five-fold in area over a period of approximately 10 minutes. Gilbert suggests that this expansion may help to explain perceptual filling-in, such as color and texture filling-in, and even illusory contours (Gilbert 1992, p. 8). The idea is that the expansion allows stimuli located near the boundary of the original receptive field to drive the cell. The cells would then fire as if the stimuli were close to their receptive field centers – leading to a shifted percept of the location of the stimulus – so that the unstimulated region would appear to fill in.

In a more recent study, DeAngelis et al. (1995a) investigated receptive field plasticity in V1 under artificial scotoma stimulation similar to that in the Pettet and Gilbert experiments. However, unlike Pettet and Gilbert, DeAngelis et al. did *not* encounter changes in receptive field size (but see Chapman & Stone 1996). Instead, they report short-term changes in responsiveness (gain changes) for some cells. However, they also conclude that the sort of receptive field changes they observed could account for psychophysical phenomena such as filling-in.

4.4. Texture filling-in and cortical cell responses

A recent study by De Weerd et al. (1995) takes a major step toward bridging the neural and perceptual levels in the case of perceptual completion. De Weerd et al. discovered cells in extrastriate cortex whose responses correlate well with

Figure 13. Schematic representation of the texture filling-in stimulus used by De Weerd et al. (1995) in their comparison of perceptual filling-in and cortical cell responses. Human subjects were asked to fixate on the fixation point (FP) and indicate when they saw the hole fill in by releasing a button. Under similar stimulation conditions, cortical cell responses were measured for cells in V2 and V3 of the primate visual cortex.


Figure 14. Responses of cortical neurons stimulated with texture pattern from the study of De Weerd et al. (1995). Schematic representation of the responses of a V3 cell in the hole condition (solid line) and no-hole condition (interrupted line). The shaded zone indicates the range of average times required for human observers to report filling-in. Adapted from De Weerd et al. (1995) with permission.

the perceptual experience of texture filling-in of the sort studied by Ramachandran and colleagues (Ramachandran et al. 1993; Ramachandran & Gregory 1991). De Weerd et al. first determined the time course of perceptual filling-in for human subjects. They used a large texture with an equiluminant hole in the middle, located 8 degrees from a fixation spot (Fig. 13). Subjects were instructed to indicate when they saw the hole fill in. As the hole size was increased from 1 to 12.8 degrees, the time required to see it fill in steadily increased. De Weerd et al. then recorded from two awake behaving rhesus monkeys that viewed the same patterns while they were rewarded for maintaining fixation. For each cell, the hole was centered over the receptive field. There were two main experimental conditions. In the hole condition, the cell responses were recorded for the texture with a hole (the same condition used for the human subjects). The other no-hole condition served as a control to establish the responses to the same texture without a hole.

Cell responses in areas V2 and V3 revealed neurons whose firing rate in the hole condition was initially lower than in the no-hole condition, but that gradually increased their responses to a similar level, exhibiting what the authors term *climbing activity* (Fig. 14). In other words, after a few seconds of fixation these extrastriate cells responded to the texture with the hole as if it were a texture without a hole. De Weerd et al. suggest, therefore, that the perceptual filling-in results from a minimization of the response differences in the hole and no-hole conditions.

4.5. Discussion

In this section, we would like to raise several questions about the Pettet and Gilbert explanation of how receptive field plasticity might subserve completion phenomena, assuming for the moment that the experimental results they obtained are valid, despite being contested by the

DeAngelis et al. (1995a) study (see Chapman & Stone, 1996, for a discussion that reconciles these two studies). (Similar proposals, though less specific, have been advanced by Fiorani et al. 1992, Ramachandran 1992a, 1992b, Churchland and Ramachandran 1993, and DeAngelis et al. 1995a).

The first issue is methodological. The proposal relies on a problematic use of the analogy family of linking propositions:

$$\phi \text{ ``looks like'' } \psi \rightarrow \phi \text{ explains } \psi.$$

In Gilbert's proposal, ϕ refers to the activities of single neurons and ψ refers to subjective report data. The two sides are then connected through a sort of "resemblance" or "analogy" – the physiological data "look like" the perceptual phenomena because they both involve kinds of completion.

Building on Teller's discussion (1980; 1984), we can raise three concerns about this sort of linking proposition when ϕ stands for the activities of single neurons:

1. What is the proposition's intended range of applicability? Applied to the present proposal: for how many of the filling-in phenomena is it supposed to hold?

2. Is there sufficient homogeneity at the neural level on which to base the proposition? Applied to the present proposal: are there many "completion neurons" and are their responses homogeneous enough to support the link to the perceptual level?

3. What about interactions between the putatively privileged set of neurons and neural activity elsewhere in the visual system? Applied to the present proposal: does activity elsewhere in the visual system (for example backprojections to V1 from other cortical areas) affect the response patterns in V1?[7]

To this third concern we can add a fourth related one. In supposing that the activity of single cells is reflected more or less directly in the psychophysically measured response, one thereby assumes that, in the highly controlled conditions of an experiment, nothing else in the system interferes with the influence those cells have on the animal's response. (Teller 1980, p. 164, calls this assumption the "nothing mucks it up proviso.") In the Fiorani et al. experiments, however, the experimental conditions involve the animals

being anesthetized. One often finds scientists making hypotheses about the physiological correlates of perception based on findings in animals that are not consciously perceiving anything due to anesthesia. (Zeki's studies of color perception and V4 are a well known case in point; Zeki 1983a; 1983b). Here the "nothing mucks it up" proviso amounts to assuming that consciousness – in the sense of being awake and alert – makes no difference to what the rest of the visual system is doing. We see no reason to believe this, and good reason not to: experiments have shown that when animals are awake and behaving in normal sensory surroundings, many kinds of neuronal response in visual cells become highly dependent on behavioral factors such as the bodily tilt of the animal (Horn & Hill 1969), the animal's posture (Abeles 1984), and auditory stimulation (Fishman & Michael 1973; Morell 1972). Moreover, studies in alert, unparalyzed monkeys reveal that both attention and the relevance of a stimulus for the performance of a behavioral task can considerably modulate the responses of visual neurons (Chelazzi et al. 1993; Haenny et al. 1988; Moran & Desimone 1985; Treue & Maunsell 1996).

The second problem is empirical and has to do with the observed time courses for filling in. There are two sorts of time courses for receptive field expansion in the Gilbert data: the first is on the order of minutes; the second is on the order of months. These physiological time courses do not match the perceptual ones, which typically seem to be on the order of seconds. For artificial scotomata, Ramachandran and Gregory (1991) report that perceptual filling-in happens on the order of 2 to 3 seconds. Similarly, Gerrits et al. (1966) reported that filling-in took several seconds for stabilized images; in contrast, Gerrits and Timmerman (1969) later reported that filling-in for patients with retinal scotomata happened "instantaneously." Finally, to add one more wrinkle, Paradiso and Nakayama (1991), in a study of the temporal dynamics of brightness perception to be discussed in section 6.2, timed the speed of brightness filling-in on the order of milliseconds – brightness signals appear to propagate at a rate of 110–150 deg/sec (6.7–9.2 msec/deg). Clearly, to establish a closer link between the perceptual data and the physiological data we must await the application of new techniques for assessing receptive field changes that might occur on the order of seconds or even shorter time scales (but see DeAngelis et al. 1995b).

The study by De Weerd et al. (1995) addresses the two issues just raised. It carefully tries to correlate the time course of perceptual completion with the cell responses investigated. The authors show that the time course of climbing activity follows that of perceptual completion as the hole size is increased, and they clearly delimit the range of applicability of their linking proposition to encompass only Ramachandran-type texture filling-in. Finally, the study employed awake, behaving monkeys under stimulation conditions that paralleled those used for the human subjects.

To conclude this section, at the present time the neural processes involved in the different sorts of perceptual completion are largely unknown. However, on the basis of the foregoing considerations, we think that it is best not to view perceptual completion as directly reducible to "atomic" neural properties at the single cell level. In other words, at the neural level perceptual completion might be better

described using concepts such as "cell assemblies" or other forms of distributed coordinated activity.[8]

5. Dennett on filling-in

We turn now to Dennett's criticism of the filling-in idea, in particular to his claim that filling-in is really just finding out.

5.1. Filling-in versus finding out

Dennett makes two kinds of points about filling-in, one conceptual and the other empirical.

The conceptual points depend on distinguishing clearly between the *content* of a representation and the *vehicle* or *medium* of representation. Suppose one sees a colored region. This is one's perceptual content. Dennett assumes that there must be states or processes in the brain that bear this very content. However, he observes that this could be accomplished by the brain in a number of different ways. First, there could be a representation of that region as colored, or a representation of that region could be absent, but the brain ignores that absence. The point here is to distinguish between the presence of a representation and ignoring the absence of a representation (Dennett 1992, p. 48). Second, suppose there is a representation of that region as colored. This, too, could be accomplished by the brain in different ways: for example, the representation could be spatially continuous or pictorial, or it could be symbolic.

These conceptual points illustrate the main mistake made by analytic isomorphism. Analytic isomorphism holds that there must be an isomorphic neural representation for each conscious perceptual content. As Dennett correctly observes, however, there need not be any isomorphism between perceptual contents and neural representations, because some perceptual contents might correspond to neural processes that ignore the absence of neural representations, or they might correspond to symbolic representations.

Take the blind spot, for example. From the fact that one has no awareness of a gap in one's visual field, it does not follow that there must be a neural representation of a gapless visual field, for the brain might simply be ignoring the absence of receptor signals at the blind spot. Nor does it follow that the blind spot must be completed with spatially continuous representations, for the region might simply be designated by a symbol.

Dennett's conceptual points still leave open the empirical matter of just what the brain does to accomplish perceptual completion. Here Dennett is not entirely clear about what he means when he says that the brain "jumps to a conclusion." In the case of the blind spot, Dennett asserts that the visual cortex has no precedent of getting information from that retinal region, and so it simply ignores the absence of signals from that area. The moral of this story is that the brain does not need to provide any representation for perceptual completion to occur; completion can be accomplished by ignoring the absence of a representation (cf. Creutzfeld 1990, p. 460). The contrast, then, is between providing a representation and jumping to a conclusion, in the sense of ignoring the absence of a representation. According to this story, the brain does not need to fill in the blind spot in the sense of providing a roughly continuous

spatial representation, nor does it need to label the blind spot region – the absence of any representation for the blind spot region is simply ignored or not noticed in subsequent visual processing. Notice that this is a case of providing content; the point is that there is no representational vehicle specifically devoted to the blind spot.

On the other hand, Dennett sometimes contrasts providing a roughly continuous spatial representation with labeling a region. In addition, he says that "filling in" means the former. Here the contrast is between, on the one hand, providing a spatial representation of each subarea within a region – filling-in – and, on the other hand, jumping to a conclusion, in the sense of attaching a label to the region all at once. In this story the brain provides both content for the blind-spot region and a representational vehicle devoted to that region, namely, a label.[9]

Dennett's slogan is: "The brain's job is not 'filling in.' The brain's job is *finding out*" (1992, p. 47). The principle of brain function being assumed here Dennett calls "the thrifty producer principle": "If no one is going to look at it, don't waste effort providing it." For example, to see a region as colored, all the brain needs to do is to arrive at the judgment that the region is colored. Whether Dennett thinks that the brain accomplishes this by ignoring the absence of a representation or by providing a label ("color by number"), he clearly thinks that filling-in the color of each subarea ("color by bit map") is not the thriftiest way to do it.[10]

5.2. Dennett's criticism and visual science

In visual science there has been a great deal of debate about neural–perceptual isomorphism in relation to filling-in, and the debates all predate Dennett's treatment (Bridgeman 1983; Grossberg 1983; Kingdom & Moulden 1989; Ratliff & Sirovich 1978; Todorović 1987; see also O'Regan 1992). In fact, in 1978, Ratliff and Sirovich argued against the need for a neural filling-in process in a way similar to Dennett. They argued that to assume that there must be neural filling-in to account for the homogeneous appearance of bounded regions is to misinterpret Mach's principle of equivalence as requiring that there be an isomorphic mapping from the form of the neural process to the form of the perceptual response. However, such an isomorphism is not logically necessary. Therefore, neither is a neural filling-in process (see also Bridgeman 1983; and Kingdom & Moulden 1989).

Ratliff and Sirovich went on to make some remarks that are interesting in relation to Dennett's discussion of Cartesian materialism:

> The neural activity which underlies appearance must reach a final stage eventually. It may well be that marked neural activity adjacent to edges [rather than neural filling-in between the edges] . . . is, at some level of the visual system, that final stage and is itself the sought-for end process. Logically nothing more is required. (1978, p. 847)

This point is similar to Dennett's that, once discriminations have been made, they do not need to be re-presented to some central consciousness system – a "Cartesian theater" (1991, p. 344). However, there is a dissimilarity as well: as Dennett's critique of Cartesian materialism and his alternative "multiple drafts" model of consciousness makes plain, the notion of a "final stage" may have no application at all. In fact, given the dense connectivity of the brain, with

reciprocal forward and backward projections, it is not clear what "final stage" could mean in any absolute sense (see sect. 8.1). For this reason, Dennett's discussion of filling-in represents an advance over Ratliff and Sirovich's.

Although neural filling-in may not be logically necessary, whether there is neural filling-in has to be an empirical question. Ratliff and Sirovich admitted this: "we cannot by any reasoning eliminate a priori some higher-order stage or filling in process. . . . But parsimony demands that any such additional stage or process be considered only if neurophysiological evidence for it should appear" (1978, p. 847). Dennett, too, admits this (1991, p. 353; 1992, pp. 42–43). What sort of evidence is there, then, for neural filling-in?

6. Evidence for neural filling-in

In section 3 we reviewed a large number of perceptual completion phenomena. We would like to draw attention to two cases here: illusory contours and the temporal dynamics of brightness and color induction. Both strike us as counterexamples to the idea that perceptual completion is accomplished by the brain's ignoring an absence.

6.1. Illusory contours

Several researchers have suggested cognitive theories of illusory contour perception, most notably Gregory (1972) and Rock and Anson (1979). In these theories, illusory contour formation is largely the result of a cognitive-like process of postulation. Illusory contours are viewed as solutions to a perceptual problem: "What is the most probable organization that accounts for the stimulus?" Although there is ample evidence for the role of cognitive influences in illusory contours, current studies point to the importance of relatively low-level processes in the formation of illusory contours.

Two lines of evidence point to an early neural mechanism for illusory contour completion: (1) neurophysiological data, and (2) psychophysical studies of the similarities between real and illusory contours.

6.1.1. Neurophysiologic evidence. As we discussed in section 4.1, von der Heydt and colleagues have shown that figures in which we see illusory contours evoke responses in a large number of cells in V2 of alert monkeys (Peterhans & von der Heydt 1989; von der Heydt et al. 1984; von der Heydt & Peterhans 1989). The cells respond as if the illusory contours were formed by real edges or lines, and they respond to variations in the figure in a way that resembles human psychophysical responses to the same variations. Although making a link between single cell activities and perceptual phenomena is problematic for the reasons reviewed in section 4.5, the evidence here seems to suggest that the perceptual completion of boundaries involves the neural completion of a presence, rather than "ignoring an absence."

6.1.2. Psychophysical evidence. Many psychophysical studies have provided evidence for a common early treatment of both real and illusory contours by the visual system (see Lesher 1995; Spillmann & Dresp 1995). For example, Smith and Over (1975; 1976; 1977; 1979) have revealed

similarities between the two types of contours in the realm of motion aftereffects, tilt aftereffects, orientation discrimination, and orientation masking.

Tilt aftereffects are particularly interesting. A tilt aftereffect will occur if one adapts for a few seconds by looking at lines oriented counterclockwise from the vertical, and then one is exposed to a test stimulus of vertical lines. The latter will appear to be tilted clockwise, away from the adapting orientation. There is compelling evidence from recent studies showing that tilt aftereffects cross over between real and illusory contours (Berkeley et al. 1994; Paradiso et al. 1989). Thus adaptation with real lines can affect the perception of illusory contour orientation and vice versa (see sect. 7.3.2).

An important question concerns the level at which real and illusory contours have similar status. Motion and tilt aftereffects are often attributed to short term habituation in early visual stages (Barlow & Hill 1963; Movshon et al. 1972). Thus the evidence from psychophysics is that real and illusory contours share internal processes at an early level of the visual system. In fact, there is considerable evidence pointing to the functional equivalence of real and illusory contours in the operation of the visual system (see Table 1 of Lesher 1995; Spillmann & Dresp 1995, p. 1347).

6.2. Temporal dynamics of brightness and color induction

There is an enormous body of literature on the spatial variables determining brightness and color induction. In contrast, there are considerably fewer studies investigating temporal variables (Boynton 1983; Kinney 1967; see Heinemann 1972). However, there are a few studies with results that speak directly to the question of evidence for filling-in.

Paradiso and Nakayama (1991) used a visual masking paradigm to investigate two issues – first, the role of edge information in determining the brightness of homogeneous regions, and second, the temporal dynamics of perceptual filling-in. They reasoned that if the filling-in process involves some form of activity-spreading, it may be possible to demonstrate its existence by interrupting it. If boundaries interrupt filling-in, what happens when new borders are introduced? Is the filling-in process affected before it is complete?

Figure 15 shows the paradigm they used as well as the basic result. The target is presented first and is followed at variable intervals by a mask. For intervals on the order of 50 to 100 msec, the brightness of the central area of the disk
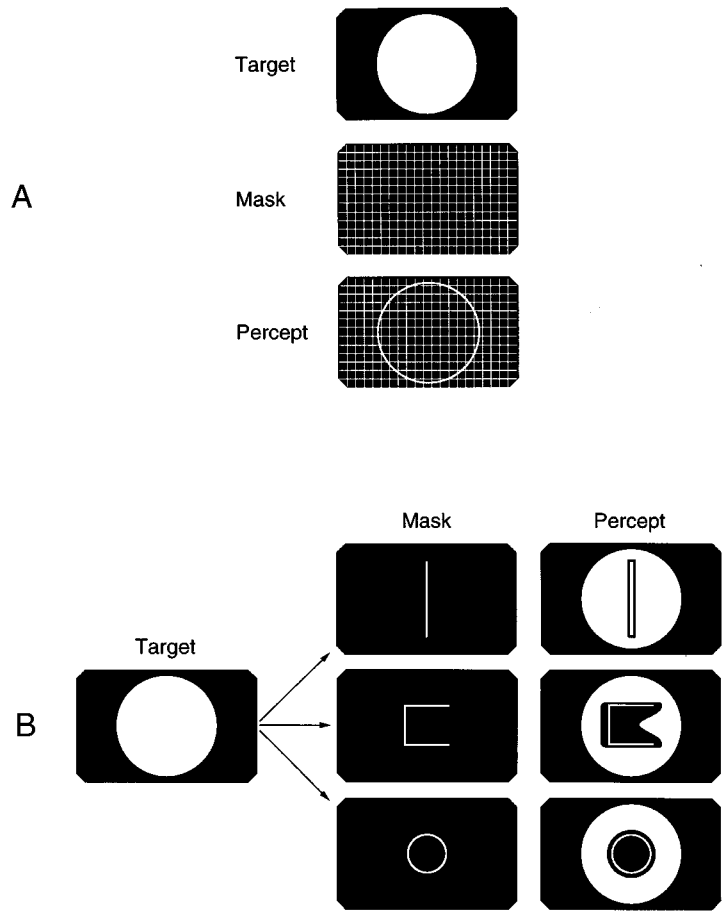
Figure 15. Masking paradigm in temporal dynamics of brightness study (Paradiso & Nakayama 1991). (A) Brightness suppression of a disk-shaped target by a mask consisting of a grid of thin lines. The target and mask are each presented for 16 msec. Optimizing the temporal delay between the stimuli yields a percept in which the brightness in a large central area of the disk is greatly suppressed. (B) Brightness suppression is highly dependent on the arrangement of contours in the mask.

is greatly reduced. If the mask is presented after 100 msec, the brightness of the central region is largely unaffected. The most striking result was that the brightness suppression depended on the distance between target and mask. In particular, for larger distances maximal suppression occurred at later times.

Paradiso and Nakayama's results are consistent with the hypothesis that brightness signals are generated at the borders of their target stimuli and propagate inward at a rate of 110–150 deg/sec (6.7–9.2 msec/deg). The idea that contours interrupt the propagation is perhaps clearest for the case in which a circular mask is introduced, resulting in a dark center, for the brightness originating from the target border seems to be "blocked." Paradiso and Nakayama discuss several alternative accounts, such as lateral inhibition processes, but do not consider them to be plausible explanations of their findings.

Some of these results were anticipated in an earlier study by Stoper and Mansfield (1978). They employed a masking paradigm in which the masks were varied systematically in time. They interpreted their "area suppression" effects as resulting from the interference of a mask with the process of filling-in of target brightness. Their paradigm enabled them to show that brightness suppression could not be due simply to contour suppression, thereby indicating that brightness and contour processes are subserved by independent systems.

The filling-in model of brightness perception proposed by Grossberg and Todorović (1988) has been shown by Arrington (1994) to produce excellent fits to the data from both Stoper and Mansfield (1978) and Paradiso and Nakayama (1991). This sort of close link between psychophysics, neurophysiology, and modeling seems especially promising for investigating the mechanisms responsible for perceptual completion.

Another relevant study comes from De Valois et al. (1986). They employed center-surround standard (reference) and matching (variable) stimuli, similar to the ones used in classic contrast studies. They compared the results of direct changes in brightness or color in which the center of the standard pattern was modulated (as was the matching pattern) to the changes that occurred when the surround was modulated sinusoidally while the center was kept constant at the mean level. These two conditions were referred to as "direct" and "induced," respectively. The purpose of the experiments was to measure the brightness and color changes produced by oscillations at various temporal frequencies between 0.5 and 8 Hz. Their studies revealed two main findings: (1) The temporal frequencies studied had little effect on the apparent brightness change in the direct condition; color variations in the direct condition were present but small. (2) In the induced condition, the amount of brightness change fell drastically as the temporal frequency increased (approximately 2.5 Hz).[11]

These results can be interpreted in terms of a spreading mechanism of induction that occurs over time, one that would provide a spatially continuous representation for filling-in. Brightness and color signals would be generated at the edges between center and surround, and would propagate inside the center region determining the appearance. An optimal temporal frequency would reflect the time interval necessary for the signal to propagate from the edges. The drastic fall-off found by De Valois et al. would result

from a change in the surround before the edge signal was able to reach the middle of the center region.

Rossi and Paradiso (1996) have replicated the brightness induction results of De Valois et al. (1986) and have studied the role of pattern size on the effect by varying the spatial frequency of the inducing pattern. The correlation found between spatial scale, degree of induction, and cutoff frequency indicates that there is a limited speed at which induction proceeds and that larger areas take more time to induce. Rossi and Paradiso conclude that the limits on the rate of induction are consistent with an active filling-in mechanism initiated at the edges and propagated inward.

In a remarkable study, Rossi et al. (1996) showed that a significant percentage of neurons in cat primary visual cortex respond in a manner that correlates with perceived brightness, rather than responding strictly to the light level in the receptive field of the cells. Rossi et al. studied cell responses in conditions analogous to the direct and induced conditions studied psychophysically by De Valois et al. (1986) and Rossi and Paradiso (1996). In the induced condition, neural responses were largest at low temporal frequencies and decreased as the rate of modulation increased over 1.0 Hz. In the direct condition, however, response amplitudes progressively increased with increasing temporal frequencies. These results, as well as other findings, are closely paralleled by psychophysical findings, suggesting that such cell responses may contribute to the perception of brightness.[12]

The studies discussed in this section provide strong evidence for featural filling-in. In brightness filling-in, the brain seems to be providing content, and it seems to be doing so through a roughly continuous propagation of signals, a process that takes time. On the other hand, ignoring a region by jumping to the conclusion that it has the same label as its surround does not take time in the same way: although labeling would involve brain processes with their own temporal limitations, there seems no reason to suppose that it would be subject to the same kinds of temporal constraints as those involved in signals having to propagate through some spatially extended area.

7. Assessing the effects of perceptual completion

Many studies of perceptual completion have provided results based directly on an observer's report of the percept. For example, observers are asked to give a verbal report of what they see in their blind spot given surrounding stimulation (e.g., Brown & Thurmond 1993). In other cases, subjects are asked to draw the shape of an occluded figure (Moravec & Beck 1986; Takeichi et al. 1995). Some studies simply ask subjects to indicate whether they perceive a given completion (as in Watanabe & Cavanagh 1993). Although many of these studies provide important information about the type of perceptual completion, how they are to be interpreted in relation to neural processes is unclear.

On the other hand, many experiments have probed perceptual completion by directly assessing the effects of the completion processes themselves. By investigating whether there are measurable effects of completion, it becomes possible to evaluate more precisely the mechanisms involved. Dennett (1993, p. 208) raises this point in a useful way: "The way to test my hypothesis that the brain does not bother filling in the 'evidence' for its conclusion is to see if

there are any effects that depend on the brain's having represented the *step,* rather than just the *conclusion.* . . . The detail would not just *seem* to be there; it would have to be there to explain some effect." In fact, visual scientists have investigated this issue of whether there are measurable effects of filling-in, as the following discussion of some relevant studies will demonstrate.

7.1. The blind spot

The motion aftereffect consists of the perception of motion when one views, for example, a screen containing stationary dots, after being exposed to motion (of the opposite direction) during a previous adaptation phase. Murakami (1995) studied the motion aftereffect after monocular adaptation to filled-in motion at the blind spot. Does the region of the blind spot (which contains no photoreceptors and so is not stimulated) also generate an aftereffect?

Instead of directly assessing whether a regular aftereffect is produced, Murakami assessed the interocular transfer of the effect, that is, whether the motion aftereffect could be measured at the corresponding visual field of the other eye. It is well known that a standard motion aftereffect transfers interocularly. Murakami found that the aftereffect also transfers interocularly in the blind spot case. In other words, adaptation to filled-in motion at the blind spot of one eye can cause a motion aftereffect at the corresponding visual field of the other eye. This result provides evidence for the perception of real motion and the perception of filled-in motion sharing a common neural pathway in an early stage of the visual system (see also note 8). If the brain treated perceptually completed motion at the blind spot and real motion differently, then one would not expect the motion aftereffect to transfer. Murakami's study thus provides a measurable effect of what appears to be the brain's having taken the trouble to fill in the motion at the blind spot, though not necessarily in a topographic manner.

On the other hand, in an earlier study, Cumming and Friend (1980) compared the strength of the tilt aftereffect induced by (partial) gratings completed across the blind spot with control gratings. Gratings were seen as completed across the blind spot, but the magnitude of the tilt aftereffect they induced suggested that the perceptually completed portions of the gratings did *not* contribute to the aftereffect. However, this negative result must be interpreted with care. One cannot rule out the possibility of completion contributing to the effect because the mechanisms involved may be at a "higher" processing level than the ones involved in the effect being probed (the tilt aftereffect).

7.2. Stabilized images

Several studies have shown that the changes in perceived color associated with color filling-in due to stabilized retinal images produce indirect effects that can be measured psychophysically. Here we briefly review some recent studies.

The sensitivity to a small-field, flickering, blue test light is significantly altered by adaptation to yellow light. Piantanida (1985) studied whether yellow adaptation induced by the filling-in of a stabilized image would have the same effect on S-cone flicker sensitivity when compared with an actual yellow light illuminating the retina. His results

showed that a yellow background induced by the filling-in of a stabilized image is as effective in reducing flicker sensitivity as an actual yellow background applied directly to the retina.

According to current thinking, flicker sensitivity can be affected by attenuation of the signals in the corresponding pathways. For example, a yellow background is known to reduce the flicker sensitivity of the S-cone system. Pugh and Mollon (1979) have proposed a two-stage model, in which attenuation would occur at the S cones themselves and at a site where S-cone signals interact antagonistically with other classes of cones. The precise mechanisms posited to account for these effects need not concern us here, but it is clear that an actual yellow light modifies the state of possibly several stages in the visual system. At the same time, Piantanida's results show that the perception of yellow in a region that is physically dark produces equivalent results. His results suggest that the perception of yellow due to perceptual filling-in is associated with the same types of changes that occur when the visual system is presented with an actual yellow light.

Nerger et al. (1993) investigated how hue cancellation and increment thresholds are affected for backgrounds that are retinally stabilized compared with nonstabilized backgrounds. They employed a disk-annulus stimulus in which the outer border was always unstabilized, and the disk could be either stabilized or not stabilized. The disk and annulus were illuminated by either 575 or 640 nm light; when one was 575 nm, the other was 640 nm. In the nonstabilized condition, subjects saw, for example, a red disk surrounded by a yellow annulus. In the stabilized condition, the color of the 640 nm disk changed from red to yellow. When subjects were asked to perform a hue cancellation of the test probe (so that it appeared neither reddish nor greenish), the stabilized and nonstabilized conditions produced different settings. (Note, however, that in both conditions the same 640 nm disk was imaged on the retina.) This result shows that the color appearance of the background can influence the color of lights superimposed on it. Nerger et al. also evaluated increment thresholds (in which an incremental test flash must be detected) for both conditions; they showed that these did not differ.[13] Thus increment thresholds are unaffected by the appearance of the adapting field and depend only on its spectral energy distribution. Nerger et al. propose a two-site model, in which filling-in affects the second (color-opponent) adaptation site, but not the first site where gain changes occur.

7.3. Illusory figures

7.3.1. Featural completion. Dresp (1992) studied whether threshold elevation (for the detection of a small light spot) would be similar throughout an illusory Kanizsa square. She reasoned that if filling-in mechanisms were responsible for producing uniform brightness levels throughout the figure, thresholds should be similar at the center of a square and near inducing elements. The results instead showed that thresholds decreased at the center of the figure.

Although Dresp interprets her results as being at odds with a brightness filling-in mechanism responsible for illusory figure brightness, there are several critical points that have to be considered. "Increment sensitivity and subjective brightness are not necessarily related by any simple

function" (Fiorentini, 1972, p. 195; see also Cornsweet & Teller 1965). For example, increment thresholds follow (more or less) the brightness of light Mach bands (Békéséy 1968), but not of dark Mach bands (see also Burkhardt 1966). Although a positive finding – constant threshold elevation throughout – would provide some indication of the levels within the visual system in which contrast elevation and brightness processes are operative, a negative finding such as Dresp's is hard to interpret because threshold elevation might be probing rather early processes.

7.3.2. Boundary completion. Paradiso et al. (1989) studied whether adaptation to illusory contours produces tilt aftereffects comparable to those obtained for regular real lines. They initially established that illusory contours used in both adaptation and test phases produce strong tilt aftereffects. Can adaptation to illusory contours induce an aftereffect when real lines are used in the test phase (or vice versa)? Paradiso and colleagues showed that the answer is yes. Adaptation to real lines induces a strong aftereffect when testing with illusory contours, but a significantly weaker aftereffect is obtained when adaptation to illusory contours is used and real lines are tested. The authors attribute this asymmetry to the corresponding asymmetry in the distribution of receptive field types in areas V1 and V2 (cells responding to illusory contour stimuli are typically found only in V2).

In another experiment, Paradiso et al. (1989) evaluated the degree of interocular transfer of the tilt aftereffect, which was found to be stronger when the test stimulus was illusory than when it was real. As the authors observe, these results are consistent with the idea that neural mechanisms activated by illusory contours are more binocular than those activated by real lines (assuming that cortical binocularity underlies interocular transfer).

In summary, the existence of a tilt aftereffect with illusory contours and its dependence on adaptation angle indicate the existence of orientation-selective neurons that respond to illusory contour stimuli. Moreover, the interocular transfer shows that real and illusory contours share an early visual pathway.

7.4. Chromatic diffusion and the gap effect

When two equiluminant colored fields abut, no clearly visible contour is seen. This occurs even for color differences of twice threshold (Eskew & Boynton 1987). This stimulus configuration can sometimes produce chromatic diffusion (Eskew 1989; Eskew & Boynton 1987). For small fields differentially exciting the S cones, the violet and green colors "bleed" across the (invisible) contour, thereby producing a larger uniform area after an initially apparent color difference sinks below threshold.

A remarkable property of the juxtaposed color-patch stimulus configuration is that color discrimination may be severely impaired. When the two fields are separated slightly discrimination improves substantially. Boynton et al. (1977) named this effect the *gap effect*. According to Boynton et al., the gap effect is related to a spatial averaging mechanism that integrates the two patches together. The small border or gap prevents such integration and improves sensitivity.

Eskew (1989, p. 717) suggested that chromatic diffusion

(the perceptual bleeding of colors across an invisible contour) was related to a physical process of integration: "chromatic diffusion . . . seemed as if it could be the visible appearance of such an integrating process, observed in real time." He determined the chromatic discrimination thresholds for juxtaposed fields as a function of stimulus duration – that is, the fields were flashed for a certain time. Discriminations were maximal at 400 msec and declined linearly (on a log scale). Eskew interpreted the approximately exponential time-course of the decrease in sensitivity in terms of a diffusive mechanism. The optimal presentation time was linked to the lack of time for the integration process to reduce the color differences across the border. Note that for such exposure durations the introduction of the gap has little or no effect. This is not the case for longer exposure durations, where the introduction of the gap improves discrimination. Also consistent with a spatial diffusive mechanism is the finding by Eskew and Boynton (1987), in which the change in sensitivity as a function of time is reduced for a short wide stimulus when compared with a tall narrow one.[14]

7.5. Summary

Experimental investigations of the blind spot, stabilized images, illusory figures, and chromatic diffusion and the gap effect provide suggestive evidence about certain consequences of perceptual completion. In particular, they point to measurable effects that seem to depend on representing a presence rather than ignoring an absence.

8. Filling-in without analytic isomorphism and Cartesian materialism

We now return to the conceptual issues surrounding the neural–perceptual relation. We have seen that there is considerable evidence for neural filling-in. The main point of this section is to show that the existence of neural filling-in does not entail either analytic isomorphism or Cartesian materialism.

Discussions of neural filling-in have been closely tied to the doctrine of analytic isomorphism. Visual scientists sometimes interpret the evidence for neural filling-in within the framework of analytic isomorphism (see sect. 8.1). On the other hand, Dennett rejects Cartesian materialism and with it neural filling-in: although he appears to concede that neural filling-in is an "empirical possibility," and says that he does not wish "to prejudge the question" (1991, p. 353), he nevertheless asserts that the "idea of *filling in* . . . is a dead giveaway of vestigial Cartesian materialism" (p. 344).

We agree that any argument for neural filling-in based on Cartesian materialism should be rejected. However, the *empirical case* for neural filling-in as reviewed above can be separated from Cartesian materialism. Hence theories and models in visual science that appeal to neural filling-in on the basis of such evidence need not be motivated by Cartesian materialism. One must distinguish sharply between the existence of neural filling-in as an empirical matter and Cartesian materialist interpretations of filling-in. Visual scientists are mistaken when they interpret the evidence for neural filling-in within the framework of analytic isomor-

phism, but it is equally mistaken to say that talk of filling-in must mean a commitment to Cartesian materialism.

8.1. Isomorphism and the bridge locus

As we discussed in section 2, the term *isomorphism* first gained prominence in visual science through the work of Köhler. Although the concept of isomorphism has often been interpreted to mean a spatial or topographic correspondence, Köhler held that neural–perceptual isomorphism should be thought of as topological or functional. Our view is that there is nothing conceptually wrong with these sorts of isomorphism as such. Whether there are either spatial/topographic or topological/functional neural–perceptual isomorphisms in any given case is an empirical question for cognitive neuroscience to decide.

What we find problematic is the doctrine of analytic isomorphism, which holds that cognitive neuroscientific explanation requires the postulation of a "final stage" in the brain – a bridge locus – in which there is an isomorphism between neural activity and how things seem to the subject. There are two critical points to be made here, one concerning the role played by the concept of the bridge locus and the other concerning the concept of isomorphism.

In their framework for mapping between the neural and the perceptual domains, Teller and Pugh (1983) call the neural structure that "forms the immediate substrate of visual perception" the bridge locus. They write:

> Most visual scientists probably believe that there exists a set of neurons with visual system input, whose activities form the immediate substrate of visual perception. We single out this one particular neural stage, with a name: *bridge locus.* The occurence of a particular activity pattern in these bridge locus neurons is necessary for the occurence of a particular perceptual state; neural activity elsewhere in the visual system is not necessary. The physical location of these neurons in the brain is of course unknown. However, we feel that most visual scientists would agree that they are certainly not in the retina. For if one could set up conditions for properly stimulating them in the absence of the retina, the correlated perceptual state would presumably occur. (p. 581)

This passage expresses a number of different ideas that need to be disentangled. First, Teller and Pugh state explicitly that a particular pattern of activity at the bridge locus is necessary for the occurrence of a particular perceptual state. However, at the end of the passage they also explicitly state that retinal stimulation is probably not necessary (assuming one could stimulate the bridge locus neurons directly), thereby suggesting that the bridge locus activity pattern is sufficient for the perceptual state. Therefore, it seems that part of what they mean by "bridge locus" is a particular set of neurons having a particular pattern of activity that is necessary and sufficient for a particular perceptual state. Second, in calling the bridge locus a particular "neural stage," and in saying that this stage is not likely to be found in the retina, Teller and Pugh seem to be conceiving of the bridge locus in a localizationist manner as a particular cortical region or area.

Analytic isomorphism relies on the concept of the bridge locus. Consider the following statement by Todorović (1987, p. 549): "A logical consequence of the isomorphistic approach is that a neural activity distribution not isomorphic with the percept cannot be its ultimate neural founda-

tion." By "ultimate neural foundation" Todorović indicates that he means the bridge locus. The doctrine of analytic isomorphism states that it is a condition on the adequacy of cognitive neuroscientific explanation that there be an ultimate neural foundation in which an isomorphism obtains between neural activity and the subject's experience.

We are suspicious of this notion of the bridge locus. Why must there be one particular neural stage whose activity forms the immediate substrate of visual perception? Such a neural stage is not logically necessary; moreover – to borrow Ratliff and Sirovich's point about neural filling-in – parsimony demands that any such stage be considered only if neurophysiologic evidence for it should appear. On this score, however, the evidence to date does not seem to favor the idea. First, brain regions are not independent stages or modules; they interact reciprocally due to dense forward and backward projections, as well as reciprocal cross-connections (Zeki & Shipp 1988). There is ample evidence from neuroanatomy, neurophysiology, and psychophysics of the highly interactive, context-dependent nature of visual processing (DeYoe & Van Essen 1995). Second, cells in visual areas are not mere "feature detectors," for they are sensitive to many sorts of attributes (Martin 1988; Schiller 1995). One of the main ideas to emerge from neuroscience in recent years is that the brain relies on distributed networks that transiently coordinate their activities (Singer 1995; Vaadia et al. 1995), rather than centralized representations. Finally, Dennett and Kinsbourne (1992) have argued that the notion of a single neural stage for consciousness hinders our ability to make sense of neural and psychophysical data about temporal perception.

Some of these critical points could perhaps be met by relying on a less localizationist conception of the bridge locus, which, as Todorović (1987, p. 550) observes, is probably an "oversimplified notion," because "there is no compelling reason to believe that the bridge locus is confined to neurons of a single type within a single cortical area."[15] Although this is a step in the right direction, the term *bridge locus* – defined as "the *location* [our emphasis] at which the closest associations between γ [physiological] and φ [psychological] states occur" (Teller & Pugh 1983, p. 588) – does not strike us as particularly useful for thinking about the distributed neural correlates of perceptual experience. For example, such correlates might involve neural assemblies in which membership is defined through a temporal code, such as response synchronization (Singer 1995; Varela 1995). For this reason, we think that the concept of the bridge locus should be abandoned.

To abandon the concept of the bridge locus means rejecting analytic isomorphism, because analytic isomorphism depends on this concept. However, some visual scientists reject analytic isomorphism while nevertheless adhering to the concept of the bridge locus. For example, Ratliff and Sirovich (1978) denied analytic isomorphism, but asserted that the neural processes involved in perception "must reach a final stage eventually." The notion of a "final stage" seems equivalent to the notion of the bridge locus. We would reject any framework that depends on the concept of the bridge locus, whether isomorphic or nonisomorphic.

We now return to the concept of isomorphism. A good example of what we object to in analytic isomorphism can be found in a statement made by Todorović (1987) in his

discussion of "isomorphistic" versus "nonisomorphistic" theories of the Craik-O'Brien-Cornsweet effect. Todorović admits that any mapping from neural to perceptual states "is an aspect of the notorious mind-body problem," but then goes on to say that "conceptually the idea of an isomorphism between certain aspects of neural activity and certain aspects of percepts may be more acceptable [than a nonisomorphic mapping], at least within a general reductive stance that assumes that, at some level of description, perceptual states *are* neural states" (1987, p. 550). We disagree. On the one hand, as Todorović recognizes, and as Köhler himself observed more than 30 years ago (Köhler 1960, pp. 80–81), the thesis of neural–perceptual isomorphism does not logically entail mind–brain identity. On the other hand, suppose one does assume that "at some level of description, perceptual states *are* neural states." Still, neural–perceptual analytic isomorphism would be plausible only if perceptual states are strictly identical to neural states (so that each type of perceptual state is identical to a particular type of neural state). However, isomorphism would not be plausible if the identity is weak, that is, if perceptual states are multiply realizable with respect to neural states (so that, although every perceptual state is identical to some neural state, one and the same type of perceptual state can be realized in many different types of neural states, or in many different types of nonneural physical states for that matter). This issue of strong (or type) identity versus weak (or token) identity is indeed "an aspect of the notorious mind-body problem," and nothing that Todorović says favors the strong identity thesis. Hence no basis has been given for the a priori claim that isomorphism is conceptually preferable to nonisomorphism in cognitive neuroscientific explanation.

9. Filling-in and the personal/subpersonal distinction

The final matters we wish to discuss are open-ended and programmatic, for they concern some of the broad conceptual and methodological issues raised by our discussion of filling-in. We have seen that arguments for filling-in based on either analytic isomorphism or Cartesian materialism must be rejected. During the course of our discussion, a fundamental conceptual point emerged, namely, that one cannot infer anything about the nature of the neural representational medium of visual perception from the character of the subject's perceptual content (see sect. 5.1). For example, suppose one has a perceptual experience that there is something red in front of one. It does not follow on logical, conceptual, or methodological grounds that there is a spatial or pictorial representation of the red region in one's brain.

We think that the full significance of this conceptual point has to do with an important distinction – the distinction between the personal and the subpersonal (this terminology comes from Dennett 1978, pp. 153–54). One must distinguish between attributions of content to the person or animal and attributions of content to the brain or nervous system (McDowell 1994). Personal level attributions treat the animal as an embodied whole embedded in an environment and as constrained by norms of rationality. In contrast, attributions of content to the brain (e.g., the visual

system) involve hypotheses about the animal's internal functional organization. In this section, we wish to show the relevance of this distinction to the filling-in controversy in visual science.

9.1. The personal/subpersonal distinction and task-level conceptions of vision

During the past two decades there has been considerable research into the subpersonal mechanisms of visual perception. One prominent research program, based on the work of Marr (1982) and Poggio et al. (1985), conceives of vision as a kind of "inverse optics" – a process of producing representations in the brain of the three-dimensional layout of objects from the limited information encoded in the two-dimensional retinal image. The central idea of this approach is that the visual system must construct an accurate representation of the world on the basis of the limited information available to the retina. However, different, nonrepresentational lines of research also have emerged in the past two decades. In particular, the "ecological approach" of Gibson (1979) and his followers (Turvey et al. 1981), as well as the more recent "animate vision" approach (Ballard 1991; 1996; Ballard et al. 1997), emphasizes not the information available to the retina, but rather the information available to the animal as it explores its environment.

We think that the distinction between the personal and the subpersonal has a direct bearing on the debate between representational and nonrepresentational approaches to visual perception (McDowell 1994; Noë 1995; Thompson 1995, pp. 232–42), and in turn on the filling-in controversy. Because the representational approach holds that vision comprises a set of complex information-processing tasks, it concentrates on representational/computational processes underlying our perceptual capabilities. These processes are all subpersonal, occurring within the animal's brain. In contrast, Gibson's ecological approach aims to provide an account, not of what goes on inside the animal, but rather of what the active, probing animal itself accomplishes in its environment. As Gibson put it: "In my theory, perception is not supposed to occur in the brain but to arise in the retino-neuro-muscular system as an activity of the whole system" (1972, p. 217). "Perceiving is an achievement of the individual, not an experience in the theatre of consciousness" (1979, p. 239). The central point made here is clear: the proper subject of perception is not the brain, but rather the whole embodied animal interacting with its environment. We believe that this point can be accepted even by those who reject the details of Gibson's specific hypotheses.

The subpersonal level is important if we wish to understand the neural mechanisms and processes that underlie our perceptual capabilities. However, the subpersonal level has influenced visual science to such an extent that the perceptual subject – the person or animal – has been neglected (with notable exceptions such as Gibson). We find this neglect unacceptable. Attention to the animal as the subject of perception is important for two interconnected reasons: first, it corrects certain conceptual problems that often emerge in the subpersonal, representational understanding of vision; and second, it suggests a better kind of task-level analysis of vision than that found in the representational approach.

Most computational and neural network models of vision

are "image-based" in the sense that they follow Marr's idea that "vision is the *process* of discovering from images what is present in the world, and where it is" (Marr 1982, p. 2). The images are patterns of light on the retinal array, and to represent what is present in the world and where it is, the content contained in the images must be extracted and reconstructed through complex internal processing. This account of vision is subpersonal because the animal – the perceptual subject – has no place in it. The problem with such subpersonal answers to the question "what is vision?" is that they lead to conceptual confusions and thus to an unsatisfactory task-level account of vision. For example, consider another statement from Marr: "The purpose of these representations [the primal sketch and the $2\frac{1}{2}$-dimensional sketch] is to provide useful descriptions of aspects of the real world" (1982, p. 43). Who is reading the descriptions? Such an approach seems guilty of the fallacy of supposing that there is a homunculus in the head whose job is to view the incoming information (Thompson 1995, pp. 234–35; for further discussion see Noë 1995). The animal, on the other hand, simply sees aspects of the world. At the level of the animal, there are no images, representations, or descriptions in visual perception (except, of course, when viewing something in the world that is an image, representation, or description); there is rather a perception-action system that enables the animal to visually guide its activity and thereby visually explore its environment. Thus by attending carefully to the level of the animal – the personal level – we arrive at a task-level conception of vision different from the representational one: the task of vision is not to produce representations from images, but rather to discover through a perceptual system what is present in the world and where it is (McClamrock 1995; Thompson 1995).

Although this kind of task-level conception of visual perception derives from Gibson (1979), it is clearly evident in other recent cognitive science research programs, such as "active" and "animate" vision (Aloimonos et al. 1988; Bajcsy 1988; Ballard 1991; 1996; Ballard et al. 1997), embodied artificial intelligence (Brooks 1991), autonomous systems (Varela & Bourgine 1992), and enactive perception and cognition (Clark 1996; Thompson et al. 1992; Varela et al. 1991). The main idea held in common by these research programs is that proper task-level analyses of perception and cognition are "activity-based" (Brooks 1991) – for example, the task of vision is to guide activity or behavior (such as hand–eye coordination in the manipulation of objects), rather than to construct an elaborate internal model of a scene. The need for representations is minimized through reliance on the perceptually guided action of the animal or system as a whole. As Brooks (1991, p. 139) observes, it is "better to use the world as its own model" than to suppose the world has to be represented in the head. O'Regan (1992, p. 484) expresses the same idea when he suggests that the environment provides an "external memory" for the animal to probe as the need arises.

What is the relevance of the personal/subpersonal distinction to the filling-in controversy? To a large extent, invocations of filling-in as a theoretical category, especially those based on either analytic isomorphism or Cartesian materialism, depend on the subpersonal, representational conception of the task of vision. Indeed, on this conception, filling-in provides a paradigm of the kind of construction on which vision depends: the job of filling-in is to complete images or representations in the brain. For example, Gross-

berg (1987a, p. 93) writes: "The images that reach the retina can be occluded and segmented by the veins in several places. Somehow, broken retinal contours need to be completed and occluded retinal color and brightness signals need to be filled in. Holes in the retina, such as the blind spot or certain scotomas, are also not visually perceived . . . due to a combination of boundary completion and filling-in processes . . . ". In addition, in another article with Mingolla: "Without featural filling-in, we would perceive a world of colored edges, instead of a world of extended forms" (Grossberg & Mingolla 1985, p. 175). However, to reject the representational conception of vision in favor of an animal-centered and activity-based conception is to downgrade the importance of filling-in as a theoretical category in the explanation of vision. This point reinforces from a different angle our earlier points that propositions about filling-in as a whole need careful consideration and that evidence for neural filling-in must be evaluated on a case-by-case basis. As we argued earlier (sects. 6 and 7), we believe that in particular cases there *is* evidence for neural filling-in, but our present point is that, once shorn of its connections to the representational conception of vision, such filling-in seems a shadow of its former self.

9.2. Recovering the personal in visual science

In the previous section, we argued for the importance of the distinction between the personal and the subpersonal in the understanding of vision. In this final section, we wish to discuss a particular assumption about perceptual content that plays a role in certain criticisms of filling-in (Dennett 1991; O'Regan 1992) and that results from neglecting the personal level. The assumption is that there is no difference in kind between perceptual content at the personal level and neural content at the subpersonal level. We reject this uniformity of content thesis. We hold that there is a difference in kind between the content of visual perception at the personal level and the content of neural states at the subpersonal level: perceptual content pertains to the animal as a whole interacting with its environment and requires for its description an animal-centered task-level account of vision, whereas neural content pertains to the animal's internal functional organization and requires for its description levels of explanation concerned with internal processing.

To see the uniformity thesis at work we need to consider some examples. Our first example is taken from Dennett's discussion of filling-in (see sect. 5.1); this example will be supplemented by two others from visual science.

Suppose someone walks into a room where the walls are covered with wallpaper whose pattern is a regular array of hundreds of identical images of Marilyn Monroe (Dennett 1991, pp. 354–55; 1992). The person would report seeing that the wall is covered with hundreds of identical Marilyns. However, the person can foveate only a few Marilyns at a time and the resolution of parafoveal vision is not good enough to discriminate between Marilyns and colored shapes. One can conclude that the brain represents that there are hundreds of identical Marilyns, but not that there is a spatial or pictorial representation of each identical Marilyn (see O'Regan 1992, pp. 474–75, 481). Conceptually, this example is analogous to the filling-in cases discussed earlier: just as the experience of a gapless visual field does not entail neural filling-in of the blind spot, so seeing that the wall is covered with hundreds of identical Marilyns does not

entail a neural representation of each individual Marilyn. Indeed, in putting forth this example, Dennett conjectures that the brain does not bother to fill in the Marilyns, in the sense of propagating a high-resolution, foveated Marilyn image "across an internal mapping of an expanse of wall"; rather, the brain just "jumps to the conclusion that the rest are Marilyns, and labels the whole region 'more Marilyns' without any further rendering of Marilyns at all" (Dennett 1991, p. 355). Yet he goes on to say: "it does not seem that way to you. It seems to you as if you are actually seeing hundreds of identical Marilyns." The implication is that in some sense the person's experience of the Marilyns is mistaken or illusory, and the reason seems to be that there is no picture in the person's brain that represents each Marilyn distinctly. The blind spot is treated in a similar way. Here too the hypothesis is that the brain jumps to a conclusion, but again "it certainly does not seem that way from the 'first-person point of view'" (Dennett 1992, p. 47). The person's experience of the blind spot being filled in is an illusion because there is no picture in the brain being filled in (the brain is really just jumping to a conclusion in the senses discussed in sect. 5.1). In general, the moral is supposed to be that although our field of view seems to be full of detail, the detail is actually an illusion (see also pp. 366, 408, 467–68). It is this conclusion and the reasoning behind it that depend, we think, on the uniformity thesis.

There are two problematic steps in the reasoning, both of which depend on the uniformity of content thesis. First, it is assumed that in the absence of a brain-level pictorial representation of each of the identical Marilyns the person cannot have a percept with the content that there are hundreds of identical Marilyns. What is striking about this reasoning is that it relies on analytic isomorphism: the underlying assumption is that it is a necessary condition of a person's having an experience that there be states in the person's brain isomorphic to how things are represented as being (for similar criticism, see Sedivy 1995, p. 475). Such analytic isomorphism is also evident in Dennett's idea that for consciousness to be really continuous, the subpersonal neural processes would have to be continuous, but they are not, so the continuity of consciousness is an illusion: "One of the most striking features of consciousness is its *discon*tinuity. Another is its *apparent* continuity. One makes a big mistake if one attempts to explain its apparent continuity by describing the brain as 'filling in' the gaps" (Dennett 1992, p. 48). In Dennett's case, the analytic isomorphism appears to be driven by the uniformity thesis – by the idea that perceptual content at the personal level *just is* the content of brain states at the subpersonal level. We accept the general thesis that facts about brain-level content determine what the person sees or experiences, but we deny that the general thesis entails the uniformity thesis and that there must be an isomorphic neural representation of the perceptual content.

The second problem comes from making the following assumption about perceptual experience: in having an experience of, for example, hundreds of identical Marilyns on the wall, it seems to one that the Marilyns are all there in one's mind or brain. Thus Dennett says of someone who claims to see all the Marilyns: "The hundreds of Marilyns in the wallpaper seem to be present in your experience, seem to be in your mind, not just on the wall. . . . But why should your brain bother importing all those Marilyns in the first place?" (Dennett 1991, pp. 359–60). Once again, the con-

clusion being drawn is that the person's experience of the Marilyns is mistaken, and the reason given is that there is no picture in the brain that represents each Marilyn distinctly. The reasoning depends on the assumption that it seems to the person that there is such a picture in his or her mind or brain. Put more explicitly: the reasoning depends on the idea that visual experience is pictorial, in the sense that to have a visual experience that is really of hundreds of identical Marilyns is to have a picture in the mind or brain with precisely that content. Clearly, to think of visual experience as being pictorial in this way depends on the uniformity of content thesis.

The assumption of the pictorial nature of visual experience is widespread in visual science, as we have seen in considering the analytic isomorphism argument for filling-in (see Todorović 1987, and our discussion of the Craik-O'Brien-Cornsweet effect in sect. 1). However, it also plays a role in the interpretations given to some experimental studies by researchers critical of filling-in. For example, building on Dennett's discussion, Blakemore et al. (1995) investigated the ability to register changes in visual scenes across saccadic eye movement. In the first experiment, they compared cases in which the image changed (or did not change) and moved in an unpredictable direction (forcing a saccadic eye movement) with cases in which the image stayed in the same place and changed (or did not change). The image changes involved the appearance, disappearance, or rotation of an object in the scene. Blakemore et al. found that when the image did not move, subjects reliably detected the changes, but when the image moved their performance fell to chance. In the second experiment, they compared cases in which the image changed and moved (as in the first experiment) with cases in which the image changed but stayed in the same place, and a mid-gray interstimulus interval separated the two images in time. This "gray-out" condition was designed to mimic what happens during a saccade. In the gray-out condition the subjects' performance was considerably reduced, though not to chance levels. Blakemore et al. interpret their results as showing "the fragility of [transsaccadic] visual memory for a complex scene" (p. 1080). They write: "we believe that we see a complete, dynamic picture of a stable, uniformly detailed, and colourful world," but "[o]ur stable visual world may be constructed out of a brief retinal image and a very sketchy, higher-level representation along with a pop-out mechanism to redirect attention. The richness of our visual world is, to this extent, an illusion" (p. 1075).

O'Regan, Rensink, and colleagues have also contributed important studies on the ability to perceive changes in scenes, and have come to the same conclusion as Dennett and Blakemore et al. (O'Regan et al. 1996; Rensink et al. 1996; for a discussion of other relevant research, going back to the 1970s, see Grimes 1996). In one study (Rensink et al. 1996), an image of a natural scene was continually alternated with a modified image, with a blank field inserted between each display. The duration of each image was 240 msec; the blank field lasted 80 msec. The modified image was the same as the original except for one change that involved either the removal of an object present in the original scene or a change in the color or spatial position of an object. Subjects found the changes very difficult to notice under these "flicker conditions," even though the changes were large and easily observable under normal conditions. Rensink et al. interpret these results as indicating that

attention is required to perceive change and that visual representations are accurate only for those aspects of a scene that are of greatest importance to the perceiver. On the basis of this experiment and another experiment investigating transients that affect attention (O'Regan et al. 1996), these authors conclude: "essentially only the 'center of interest' of a picture is encoded in memory. The impression of continuously seeing 'all' of a visual scene may be an illusion deriving from the fact that any change usually creates a visual transient that attracts attention to the changing location" (O'Regan et al. 1996).

We think that these studies provide a promising approach for investigating visual memory and visual attention, but we think that it is problematic to interpret them as showing that our impression of the visual world is a kind of illusion, for such an interpretation depends both on the uniformity of content thesis and the thesis that visual experience has a pictorial character.

The main problem with both theses is that they are based on neglecting the animal-centered or personal-level account of the task of vision. They result from trying to understand perceptual content at the level of what goes inside the animal, rather than at the level of what the animal accomplishes in its environment. At the level of the animal – the personal level – the task of vision is to guide activity in the world, not to produce internal representations.

For example, it is true that, in seeming to see the hundreds of identical Marilyns, it does not seem to you that your brain jumps to a conclusion. However, there is no reason to think that it should seem that your brain jumps to a conclusion. After all, suppose your brain actually filled in a detailed Marilyn image through some internal spatial map. It would not seem that way to you either. In general, it does not seem to you – the person – that your brain does anything one way or another when you perceive; you simply perceive the world (and, on occasion, misperceive it, as in the blind spot demonstrations).

This point is really a conceptual one about how to understand perceptual content. The claim is that perceptual content has to be understood in relation to task-level accounts at the personal level, because perceptual content has to do with the world, not with representations (pictorial or otherwise) in the brain. As we suggested in section 9.1, at the personal level, there are no pictures, images, or other representations in visual perception; there is simply experience of the world. For example, when viewing the blind spot demonstration in Figure 3, what seems to the subject to happen is not that the brain fills in some inner image or picture (or that the brain jumps to a conclusion), but rather that something in the world – the black spot on the page – disappears while being replaced (or perhaps covered over: see Durgin et al. 1995) by what is present in the surround. Similarly, the Marilyns do not seem to be present in one's experience or in one's mind (whatever that might mean); they seem to be present there on the wall. More generally, one does not have the impression of a complete, dynamic *picture* of a stable, uniformly detailed, and colorful world (a picture that may or may not correspond to the subpersonal neural representations); one has the impression of a stable, colorful world (and, we might add, the impression is in the main correct: the world *is* stable and colored).

The importance of this point for visual science is twofold. First, in devising experiments and in interpreting their results, care must be taken to ensure that the evaluation of perceptual content is not determined by subpersonal considerations about internal processing, but rather by task-level accounts at the personal level. Second, to recognize the importance of the personal level in this way means that, within the embracing context of cognitive science, scientific research into the subpersonal mechanisms of visual perception needs to be complemented by careful studies of perceptual content at the personal level (for further discussion, see Thompson et al., 1998).

10. Directions for further research

Our intention in this target article has been to offer a guide to perceptual completion for visual science and the philosophy of perception. In recent years, perceptual psychologists and visual neuroscientists have provided a wealth of new studies of perceptual completion, and these studies have attracted the attention of philosophers. Psychophysicists have produced a large body of experimental results, while neurophysiologists have begun probing the mechanisms to determine the neural substrates of the various sorts of perceptual completion. At the same time, theoretical modelers have tried to produce formal specifications of the mechanisms involved, while philosophers have tried to clarify some of the conceptual issues that arise in discussions of perceptual completion. In this target article, we hope to have shown that close collaboration between these four fields is indispensible for the visual science of perceptual completion and for the philosophy of perception. Without such collaboration, visual scientists run the risk of producing isolated facts that do not contribute to an integrated understanding of perceptual completion, while philosophers run the risk of ignoring important experimental and theoretical studies that bear on the fundamental conceptual issues.

The collaborative framework that we have tried to provide here suggests several specific directions for further research:

First, we encourage explicit consideration of linking propositions in both theoretical and experimental work. For example, when investigators suggest that a set of cells in a given region of the visual system is the neural substrate of a perceptual phenomenon, it should be clearly stated exactly how the presumed mapping is intended to hold. In particular, it is necessary to show not only that certain cellular properties correlate well with perceptual events, but that the lack of perceptual completion is paralleled by an absence of the underlying cell responses – as done by von der Heydt et al. (1984) and De Weerd et al. (1995).

Second, in psychophysics there has recently been considerable interest in the nature of modal and amodal completion. Are these subserved by common mechanisms or are they distinct? One promising line of investigation attempts to reveal whether certain experimental manipulations affect the two types of completion in similar ways. We suggest that a complementary line of investigation be employed to assess the effects of completion. Do modal and amodal completion produce similar types of "by-products" or not? In general, we advocate that studies of perceptual completion investigate the potential effects of completion processes in an attempt to determine more precisely the mechanisms involved.

Third, much of the evidence for neural filling-in discussed

above derives from experimental designs that presuppose the "image-based" or representational task-level conception of vision, rather than the active, animal-centered one. We believe that more "ecologically valid" studies of perceptual completion are needed, by which we mean studies based on the active, animal-centered task-level conception of vision.

Finally, we advocate careful studies of perceptual experience at the personal level as a complement to experimental research. Whenever we attempt to understand and explain our perception, we do so on the basis of our own perceptual experience of the world. Unless this experience has been conceptually clarified and systematically described, our scientific explanations will always be incomplete.

NOTES
1. Author to whom all correspondence should be addressed.
2. See for example the debate between Ramachandran (1993b; Churchland & Ramachandran 1993) and Dennett (1993; 1996a), and the exchange between Durgin et al. (1995) and Ramachandran (1995). See also Akins and Winger (1996).
3. Köhler's "field physics" concept received a definitive blow from a series of studies by Lashley, Sperry, and colleagues, demonstrating intact vision in cats in which the putative fields had been disrupted by the introduction of needles or by insulating regions of cortex from each other (Lashley et al. 1951; Sperry et al. 1955; Sperry & Miner 1955).
4. In view of the importance we place in section 9.2 on careful descriptions of perceptual experience, we would like to note here that calling such contours and figures "illusory" is problematic and reflects certain methodological and theoretical assumptions (Spillmann & Dresp 1995, p. 1341). We prefer the terms *apparent contours* and *anomalous contours* as being less theoretically biased. Nevertheless, *illusory contours* is widespread in the literature and used in the studies we cite. We therefore follow this usage.
5. Some form of the idea that both boundary and feature are involved in surface appearance has been proposed by various investigators, including the initial proponents of the "filling-in" idea (Davidson & Whiteside 1971; Fry 1948; Gerrits & Vendrik 1970; Walls 1954). For example, Gerrits and Vendrik (1970) spoke of the "spread of brightness-activity" (or feature) in a higher center being halted by a "brightness-barrier" (or boundary). Stoper and Mansfield (1978) also provided important evidence that brightness and contour processes are subserved by independent systems. Nevertheless, Grossberg and Mingolla (1985) were the first to propose the idea of independent boundary and featural *completion systems.*
6. Additional evidence for neon color spreading involving two distinct processes was obtained by Takeichi et al. (1992). Their first experiment showed that local color spreading does not occur under dichoptic presentation, suggesting that important monocular mechanisms are involved in the spreading effect. On the other hand, they showed that illusory contours, when not specified by each monocular configuration alone, can be determined by an interocularly combined configuration, suggesting the involvement of binocular processes.
7. Here we depart significantly from Teller's formulation (1980, pp. 164, 169; 1984, pp. 1241–42). Teller describes the analogy family of linking propositions in its application to single cells as involving a "nothing mucks it up proviso." By this she means that one has to assume that the response pattern in the single cells is preserved all the way to the bridge locus neurons – those neurons "whose activities form the immediate substrate of visual perception" (Teller & Pugh 1983, p. 581). We wish, however, to separate the idea that neural activity elsewhere in the cortex might affect the response patterns of the single cells being considered from the idea that the pattern has to be preserved all the way to the bridge locus. The reason is that we find problematic the notion of the bridge locus (see sect. 8.1). Accordingly, when we refer to Teller's "nothing mucks it up" proviso in the following paragraph we are referring *only* to the assumption that nothing else in the visual system (or elsewhere for that matter) interferes with the influence the single cells have on the animal's response; we are in no way endorsing the futher point that invokes the bridge locus idea.
8. Albright (1995) discusses a recent study by Assad and Maunsell (1995) of the neural correlates of inferred motion in monkey posterior parietal cortex, arguing that it provides evidence for the existence of "depictive representations" in amodal perceptual completion. Assad and Maunsell collected single cell recordings for awake behaving monkeys in three experimental conditions. On full vision trials, a stimulus appeared 12 degrees from the fovea, and after a brief delay moved toward the fixation spot; in this and other conditions the monkeys maintained fixation throughout the trial. On occlusion trials, the stimulus appeared as above, was invisible for a brief duration, and finally reappeared at 2 degrees of eccentricity moving toward the fixation spot. The disappearance time was consistent with a constant speed movement throughout the trial. On blink trials, the stimulus appeared and disappeared exactly as in the occlusion trials, but then it reappeared in the same peripheral location (as if it had been stationary during the period in which it was invisible). Assad and Maunsell found a large proportion of neurons that were more active following the disappearance of a visual stimulus in trials in which the animal could have inferred that the invisible stimulus was moving rather than stationary. They propose that the difference in activity between the occlusion and blink trials could be the neural correlate of such an inference. Contrary to Albright, we do not think that this study is sufficient to determine the nature of the underlying representations. Indeed, Assad and Maunsell are more cautious than Albright – they conclude only that they may have uncovered a neural correlate of the animal's inference about motion, not that they have determined the nature of the motion representation itself.
9. In visual science the term *filling-in* often has the sense of providing a roughly continuous spatial representation of a visual region. However, the difference between Dennett's two stories enables us to see how there could be a purely symbolic–inferential version of the filling-in debate (compare Nakayama & Shimojo 1990a). Here the issue would be: how many of the premises need to be "filled in" in inferential processing? How often can one simply "jump to the conclusion" without bothering to fill in the premises? Compare Dennett (1993, p. 208): "The difference I am after is the difference between jumping to a conclusion and stepping to a conclusion by making some bogus steps on which to rest the conclusion (e.g., paint in the region, and then use that painted-in region as one's 'evidence' for the conclusion subsequently drawn)." Notice how Dennett's statement runs together the elements we are trying to disentangle.
10. Dennett's position on filling-in has provoked a critical

response from Churchland and Ramachandran (1993). They disagree with Dennett because they think that the brain is not merely "finding out" in the perceptual completion of the blind spot and artificial scotomata. They also disagree with Dennett's claim that there are no neural responses devoted to the blind spot, citing Fiorani et al. (1992) as showing the contrary. As far as we can see, however, Churchland and Ramachandran do not think that the brain fills in in the sense of providing a roughly continuous spatial representation. In fact, on the basis of Ramachandran's other writings (1992a; 1992b; 1993a; Ramachandran & Gregory 1991), it seems to us that they might be prepared to accept some variant of the second story we attribute to Dennett – the one in which the brain attaches a label. In any case, we think the debate would be better conducted in relation to issues about isomorphism and linking propositions. We go on to discuss this in section 5.2.

11. Note that these frequencies are much lower than the frequencies usually revealed in flicker studies, which have cut-off frequencies of more than 30 Hz and peak around 4 to 6 Hz.

12. Given that the experimental animals (cats) were anesthetized, the same caveats discussed in section 4.5 apply here when interpreting the relationship between neural and perceptual events.

13. These results are consistent with the work of Cornsweet and Teller (1965), which showed that increment thresholds are unaffected by changes in the appearance of backgrounds when the physical characteristics of those backgrounds are held constant.

14. We cannot review here all the details of the experiments performed by Eskew (1989), but we would like to note that several of his findings suggested that a complete account in terms of adaptation and eye movements is not likely: "Although eye movements and adaptation might play a role in the chromatic diffusion phenomenon, an additional mechanism which is sensitive to the effect of the contour would be required to account for the gap effect" (Eskew 1989, p. 726).

15. Nevertheless, Todorović goes on to say that, "given the lack of relevant data, the single-cell type of bridge locus is a heuristically useful assumption." We disagree, for the reasons stated in section 4.5.

Open Peer Commentary

Commentary submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

Analytic isomorphism and speech perception

Irene Appelbaum

Department of Philosophy, University of Montana, Missoula, MT 59812.
appel@selway.umt.edu

Abstract: The suggestion that analytic isomorphism should be rejected applies especially to the domain of speech perception because (1) the guiding assumption that solving the lack of invariance problem is the key to explaining speech perception is a form of analytic isomorphism, and (2) after nearly half a century of research there is virtually no empirical evidence of isomorphism between perceptual experience and lower-level processing units.

A problem with some work in cognitive science is that the empirical data are too highly leveraged theoretically. Pessoa, Thompson and Noë's target article provides terrific relief from such work, for its significance lies precisely in its deflationary theoretical inter-

pretation of the empirical findings it cites. There is empirical evidence for neural filling-in, but its importance is doubly qualified. For one thing, neural filling-in cannot be promoted to a methodological principle. Empirical isomorphism does not imply analytic isomorphism. Moreover, even in cases where there is neural filling-in or isomorphism between neural activity and perceptual experience more generally, it becomes less significant once we see that perception is not (just) a matter of what the brain is doing, but of what the animal is doing as well. Put differently, if we do not care (as much) about representations, *a fortiori,* we do not care about whether they are isomorphic to the subject's experience.

Still, the central methodological moral of "finding out about filling in" is not, I think, that we should downgrade the status of filling-in as a theoretical category, but that we should reject analytic isomorphism. For as Pessoa et al. point out, analytic isomorphism states a criterion of adequacy not just for explanations of vision, but for cognitive neuroscientific explanations more generally. By way of reinforcing this claim, I would like to suggest that their criticism of analytic isomorphism is especially applicable to the domain of speech perception. Indeed, I will suggest that rejecting analytic isomorphism provides the key to progress in explaining speech perception.[1]

The particular form that analytic isomorphism takes in speech perception research is the assumption that the goal of a theory of speech perception is to solve "the lack of invariance problem." With a few recent and notable exceptions,[2] solving the lack of invariance problem has been taken to be a criterion of adequacy for explanations of speech perception for nearly fifty years. "Lack of invariance" refers to the widely recognized fact that there is no one-to-one correspondence between units of acoustic structure and perceived phonemes. A single consonant sound, for example, may be realized acoustically in a variety of ways depending on the surrounding phonetic context. This lack of invariance only becomes a problem, however, if one assumes that there must be an isomorphic mapping between the subject's perceptual experience and underlying physical structure. And this is just what has been assumed.

Thus, identifying an initial stage of processing in which processing units *are* structurally isomorphic (e.g., segmented, linearly concatenated) to perceived phonemes is taken to be a criterion of adequacy for an explanation of speech perception. Although these processing units are not always explicitly identified in neural terms (sometimes they are identified in acoustic or articulatory terms), it is the doctrine of analytic isomorphism that underwrites the goal of solving the lack of invariance problem, for the latter assumes that there must be isomorphism between the subject's perceptual experience and lower-level processing units (called "invariants").

As it is put in one well-known article, "invariant gestures[3] of some description there must be, for they are required, not merely for our particular theory of speech perception, but for *any* adequate theory [emphasis in original]" (Liberman & Mattingley 1985, p. 3). The difficulty is that speech perception research guided by the framework of analytic isomorphism, has made little progress. Put bluntly, by the standards embodied in analytic isomorphism, there are *no* adequate theories of speech perception. Despite nearly half a century of searching, virtually no empirical invariants have been found. Only one theory has identified even a candidate for an invariant property, but it has done so for only one class of consonants, and it is present only about 85% of the time (Stevens & Blumstein 1981). Thus, is seems that analytic isomorphism, as embodied in the goal of solving the lack of invariance problem, will have to be rejected before progress in explaining speech perception will be made.

Pessoa et al. caution against inferring analytic isomorphism from the evidence of particular cases of empirical isomorphism (e.g., cases of neural filling in). But in the domain of speech perception, analytic isomorphism survives despite the virtual absence of empirical confirmation. This suggests that in speech perception, at least, analytic isomorphism functions not merely as a methodological principle, but as dogma. If so, it is likely to prove

more difficult to eradicate. A personal-level, activity-based approach, however, such as the one favored by Pessoa et al. for vision, may well be the place for speech perception research to start.

NOTES
   1. In Appelbaum (1995) I treat this claim much more extensively, although I do not use the term "analytic isomorphism."
   2. These include Browman and Goldstein (1996), McClelland and Elman (1986), and Nusbaum and Henley (in press).
   3. Gestures are here taken to be forms of neural representations.

Filling-in as the phenomenal side of binding

Karl Frederick Arrington

Psychology Department, Arizona State University, Tempe, AZ 85287-1104; and Visual Research Laboratory, Raytheon Training, Inc., Mesa, AZ 85206-0904. **k.f.arrington@ieee.org   www.swlink.net/~kfa/index.html cns-web.bu.edu/pub/karl /www/home.html**

Abstract: The question is broadened from isomorphism to invertible transformation and optimal representation. Motivations are drawn from image compression but with an emphasis on object segmentation. Filling-in is considered as the phenomenal side of the binding process with back-surface filling-in being important. Finally, re-normalization of local filtering by globally integrated context is emphasized.

The question is not whether the visual system is isomorphistic, for clearly it is not by at least some cognitive level, but rather how the system transforms the information with minimal loss and optimal representation abstraction. The image is already distorted by the lens and needs some reconstruction at the regional level. Visual area one (V1) has applied a log polar transform, but we accept this as isomorphic by virtue of its being one-to-one, where one-to-one is important because of invertibility. Also we easily accept that at least some information has been lost by V1.

Now we have begun down the slippery slope: we have accepted a transformed representation and we have accepted at least a minimal loss of information. This is the same issue as in still- and moving-image compression. We see near lossless compression by 2D and 3D Gabor filters. This is one-to-one and can be completely lossless. So where we draw the line for isomorphism is not as interesting as how the visual system attempts to organize and efficiently represent 2D and 3D visual scenes and 4D visual events. It is reasonable for a scene in depth to be represented as roughly 3D log-polar at V1 or V2.

Now the question: Is motion represented isomorphistically? Clearly, even at the level of a motion detector there is an encoding of motion, but this could be one-to-one, so it does not violate our criteria. Now I like to think of the 3D log-polar scene as changing in time and also with a higher resolution at near times than farther. So we have a 4D log-polar scene that is begging for a filter encoding and representation, all of which, at least in theory, is one-to-one. These 4D filters could be grouped into 4D hypercomplex codes that we arrange by serial coding into symphonies of local 3D events in time.

Next we consider how the local junction events are grouped into objects that we can easily imagine to move as cohesive-matter entities. The most important thing is for the animal to be able to survive. There needs to be visual segmentation into objects: food, mates, obstacles, and paths. These objects must be manipulated, in the case of food, into the mouth. For some animals this means being graspable, for others, at least nudgeable. Edge junctions (e.g., T-, Y-, and Arrow-junctions) must bind to form an object that is manipulable. Filling-in can in many ways be considered as the phenomenal side of the binding process, where surfaces that bind the junctions and high curvature regions are constrained within their boundaries, like a soap bubble in a wire loop (a kind of

boundary value problem). The junctions (hypercomplex cells with local relative-depth indications) bind.

Filling-in is the phenomenal manifestation of active binding within local surface constraints, which forms a segmented object – the pieces of which will all be conceived as moving together, for example, when nudged during foraging. Moreover, the junctions specify expectations of back surfaces, though not *seen* as such, these surface expectations guide grasping and nudging behavior with great accuracy and relatively few surprises. Violations of expectation occur from *accidental 3D objects* that violate symmetry (e.g., solids of revolution) expectations. So there is a front-surface (modal and amodal) phenomenal filling-in side to binding into objects, but there is also a back-surface-imagination that is also a phenomenal filling-in side of binding

Finally, for surface segmentation and object binding to proceed correctly, local filter information must be interpreted (re-normalized) in terms of the globally integrated context. For example, it has been shown that global transitive luminance scene relations can be lost by using only local oriented-contrast and filling-in operations. But by building local feature filling-in operations upon more distal information, as is with Directional Filling-In (DFI) theory (Arrington 1996), the phenomenal surface feature segmentations are preserved.

Filling-in as a *within-level* propagation may be an illusion

Talis Bachmann

Department of Psychology, University of Portsmouth, Portsmouth PO1 2DY Hampshire, United Kingdom. **talis.bachmannt@port.ac.uk**

Abstract: "Finding out" about the visual world as approached from the organismic level may well include the "filling-in" type of perceptual completion if considered in terms of underlying neurophysiological mechanisms. But "filling in" can be interpreted not only as a result of within-level propagating of neural activity, but as a byproduct of the process that is necessary for modulating preconscious information about physically present objects or events so as to generate conscious quality in attending to them.

If two alarm clocks in two neighbouring rooms happen to ring simultaneously, this does not mean that the activity of ringing "propagates" from one clock to another. The apparent propagation may be an accidental byproduct of the fact that somebody set the clocks for the same time. This example illustrates what will be the main argument of this commentary, namely, that there is insufficient evidence in support of the neural "filling-in" as a lateral, within-level propagation of activity from the neurons that represent actually presented and directly valid sensory data to the neurons that represent contextually valid, but receptively invalid sensory data. It seems more likely that the "propagation" and interpolation processes assumed to be the basis of perceptual completion and filling-in are mediated by a "third process" that originates from a level that is different from the one where specific, cortical, sensory neurons are located.

The first argument is more general and can be explained as follows. Perceiving organisms are adaptively successful insofar as they manage to find a compromise between two important requirements: (1) They should not miss any potentially dangerous or otherwise important signals (data), and in doing so they should sustain adequate and consistent responding to the full concreteness of the environmental stimuli from as wide an array as possible and (2) They should also be capable of generalising and finding out abstract, meaningful characteristics of the environment without losing them in the ever varying concrete appearance of stimulation (e.g., variability in lightness, shading, size, angles, movement vectors, etc.) and despite deformations, "crowding," and noise. The first requirement is that of full, "merciless," *actual veridicality* but the second is that of flexible categorizability,

which is sensitive to expectations and maintains *potential typicality* of appearance, given certain contexts. It is obvious that a single level of the neural representation system is incapable of satisfying both these requirements at once because they are incompatible. To fulfill the task, more than one level of data analysis is necessary, featuring flexible, dynamic ways of mapping the interlevel activities.

If perceptual completion or filling-in were the function of some sort of propagating or spreading activity only within the unitary neural system at the sensory-representational level, then the first requirement (actual veridicality) would be sacrificed. This system would hence be unlikely to create illusory experiences in itself. If "finding out" and "ignoring the absence" (both governed by the principles of context sensitivity and typicality) were the only ways the neural levels could deal with data, then the adaptive advantages would be lost again and adequate responsiveness and authentic monitoring would be replaced by wishfulness. Pessoa et al.'s target article is useful in showing that to follow the strategy of "finding out" as the good one at the organismic level, perceptual systems sometimes produce neural filling-in (indicating the presence) at the lower level. It seems to me, however, that so as not to violate the first requirement of maximum veridicality it is quite probable that neural filling-in at a lower level has its origins, not in the process of intralevel lateral propagation, but in the input from some other level of sensory information processing. From where, then?

The second argument is related to the neurophysiology of sensory processes. A neuron tuned to a certain specific and spatially located characteristic can be activated in various ways. The main sources of activation include: (1) the input through ascending primary pathways that originates from receptors and mediates responses to the actually exposed stimuli for which the neuron is veridically tuned; (2) top-down input from the higher, context-sensitive levels; (3) lateral input from within the same level of neuronal analysis; (4) modulatory input from the nonspecific arousal system (NSP: e.g., thalamic nonspecific nuclei) that has been shown to be necessary for "upgrading" the activity of the specific neurons so that perceptual information carried by them will participate in conscious experience (see Bachmann 1997, for a review). It is clear that the origins of neural activity in neural filling-in can be related to (2), (3), and (4). Although the first argument in paragraph 2 biases me against the uncontested inclusion of (3) among the main possible neural filling-in mechanisms, it seems fair enough to say that at the present state of our knowledge it would be reasonable to regard all three as potential candidates. Whereas all phenomena discussed in the target article, however, are related to perceptual *conscious experience,* it seems parsimonious to discuss the role of (4) as the source of neuronal activity that has been *prima facie* linked to the function of consciousness in mediating the neural filling-in.

In recent years we have been presented with strong evidence that neurons in visual cortical areas continue to respond to exposed (actual) perceptual information to which they are tuned to even if this information is not directly present in the perceptual awareness of the subject (e.g., He et al. 1996; Leopold & Logothetis 1996; Logothetis et al. 1996). This means that it is insufficient to find the correlates of interpolated perceptual features as certain neural activities in visual cortices to posit them as the equivalents of perceptual filling-in. We should be able to find neural processes that are present with perceptual completion and absent when this completion is not phenomenally experienced.

My hypothesis is that in many cases the filling in is a *byproduct* of the thalamocortical modulation of the representational neurons that stand for actually present stimuli. Two thirds of cortical visual neurons have presynaptic input from the NSP (Brooks & Jung 1973). It is known that both the spatial and temporal resolution of the NSP is considerably poorer than those of the specific sensory-representational neurons. When sensory stimuli are presented, then in addition to evoking fast specific encoding, they evoke slower and spatially spreading modulation from NSP where neigh-

bouring units, the specific signals of which have not been presented, may also become modulated by this NSP process. (Modulation means enhancement of excitatory postsynaptic potentials, which increases the cumulative firing frequency of the modulated neurons.) As a result, "filling-in" can be expected. As primary sensory pathways are semiautonomous from NSP, we can expect perceptual completion even if the input that provided the context is not itself visible (e.g., as a result of masking).

For example, Gellatly (1980) showed that illusory-contoured figures can be perceived even when the "pacmen" that induce the "boundary completion" are not perceived because of metacontrast masking. The findings of Paradiso and Nakayama (1991) and of Stoper and Mansfield (1978) (sect. 6.2, paras. 2–5) can be explained as the outcome of a gradual arrival of the temporally delayed and spatially spread NSP modulation that leads to the conscious quality of the contoured stimulus *and* its immediate surround (notice that the surround quality belongs to the background of the contoured stimulus). Perhaps, just as alarm clocks ring synchronously not because of any "propagation" of ringing, but because of their common master, the neurons responsible for actual and interpolated sensory qualities fire together because they share the modulator which is located elsewhere.

Prediction, inference, and the homunculus

Horace B. Barlow

Physiological Laboratory, University of Cambridge, Cambridge CB2 3EG, England. **hbb10@cam.ac.uk**

Abstract: Prediction, like filling-in, is an example of pattern completion and both are likely to involve processes of statistical inference. Furthermore, there is no incompatibility between inference and neural filling-in, for the neural processes may be mediating the inferential processes. The usefulness of the "bridge locus" is defended, and it is also suggested that the interpersonal level needs to be included when considering subjective experience.

Pessoa et al.'s is a very comprehensive and useful review, and I want to suggest two additions. First, prediction should be considered because it is pattern completion in the temporal dimension; it certainly depends on statistical inference and shows that this is likely to be important in all completion phenomena. The second suggestion is to add the *interpersonal* to the authors' illuminating *personal/subpersonal* distinction, for this may tell us the origin of the homunculi who look at the images in our brains.

Statistical inference, prediction, and completion. Helmholtz thought that perceptions were unconscious conclusions about external causes inferred from the fragmentary and incomplete evidence provided by sensory messages (Helmholtz 1925, Ch. 26, p. 4) and this is still a leading concept in modern attempts to understand and model perception. Such conclusions depend upon statistical inference, and this is typically most useful just when filling-in and completion phenomena occur – when the evidence is incomplete. For instance, you gain little by identifying a regular sequence of events after it has occurred, but if you identify it before it has run to completion you can make a reliable prediction – with obvious advantages to survival in the jungle, on the stock market, and everywhere else. Prediction is pattern completion in the temporal domain; most examples of filling-in apply to static patterns, but in both cases the appropriate statistical work has to be done to gain reliable benefits, and it is no exaggeration to say that this work plays a role in perception as important as the conversion of chemical energy to mechanical work in muscle contraction.

Statistical inference requires a neurophysiological mechanism, and Pessoa et al. point out in their Abstract that "Certain forms of visual completion seem to involve spatially propagating neural activity (neural filling-in)." This could be the inference mechanism

in action, and it is hard to see why the authors go on to say that it is contrary to Dennett's (1991) description of filling-in as the brain's "ignoring an absence" or "jumping to a conclusion." Statistical inference, jumping to a conclusion, subjective filling-in, and spatially propagating neural activity are surely different levels of description of the same process in the brain and there need be no incompatibility between them. If one thing is certain it is that jumping to conclusions and ignoring absences will not take place in the appropriate circumstances unless there are neural mechanisms that approximately compute the statistics required, and the need for these neural mechanisms clearly places the spatially propagating activity in a new light.

Statistics and predictability also lead us to a psychophysical linking proposition that, unlike all the other ones Pessoa et al. discuss, is irrefutable: *Whenever two stimuli can be distinguished reliably, then some analysis of the physiological messages they cause . . . would enable them to be distinguished with equal or greater reliability.*[1] Since the Shannon-information in sensory messages cannot be increased as it passes through the brain, the states of activity in underlying neural representations *must* provide sufficient information for the perceptual discriminations that are based on them, and this enables one to find out when brain events can and when they cannot explain perceptual phenomena. A good feature of this linking proposition is that it moves on from the metaphysics of psychophysical links and instead focuses on the reliability of discriminations among physical stimuli, sensory messages, and perceptions, for these are open to experimental measurement. Its greatest strength, however, is that its truth cannot be denied, for it is simply the second law of thermodynamics applied to information flow in the brain: it is possible to lose information, but it is not possible to create it (Brillouin 1956).

One of the concepts this linking proposition can clarify is the "bridge locus" of Teller and Pugh (1983). This can be interpreted as the location in a neural system where one can most readily show that particular neural states provide enough information to account for particular psychological performance. Examples where such locations have been found include retinal ganglion cells that can detect small, brief light flashes at absolute threshold (Barlow et al. 1971), V1 neurons that can detect displacements in the hyperacuity range (Parker & Hawken 1985), and neurons in MT that can detect weak coherent motion in random dot kinematograms with sensitivity as great as the whole animal (Newsome et al. 1990). On the other hand, there are no known retinal ganglion cells that could detect the weak, spatially extended, long-lasting stimuli that intact subjects can detect (Denton & Pirenne 1954), and to express this we simply say that we do not know where the bridge locus for that task is, and cannot yet be sure whether there is one. Collecting the statistical evidence required to perform a particular psychophysical task is not a trivial matter (Barlow & Tripathy 1997), and it is useful to have a term to describe areas that house the neurons that do it.

The homunculus and interpersonal factors. Turn now to a feature of the target article that I whole-heartedly admire – the introduction of the subpersonal/personal distinction into the discussion of subjective phenomena such as completion. Can we not extend this to include interpersonal interactions? This extension may at last give the homunculus, that scandal of modern perceptual psychology (Barlow 1996), a respectable role.

To show how this may come about consider the Gibsonian example of a pilot landing a plane. He notices that the expansion point in the flow field in front of him is not exactly where he has been taught that it should be, and adjusts the controls to correct this while keenly observing the flow field to see whether his corrections are being effective. We tend to postulate a homunculus to do this inspection, in spite of the fact that almost everyone recognizes it as an unscientific intrusion when considering subpersonal mechanisms, and no better in this example of the pilot's personal interaction with the environment. But go further and consider the interpersonal interactions between the instructor and the pilot when he was being taught: wasn't the instructor's role

to implant into the pilot's brain the homunculus who inspects the flow-field? This single example does not change our thinking greatly, but we depend on homunculi more than we like to admit, and it makes a considerable difference if *most* of them are agents, implanted in our minds by others, who enable us to see the world as they see it. This would give homunculi great importance in our social lives, and would suggest one reason why our conscious experiences can quite easily be communicated to others. Although it makes our own experience dependent on others in a way that many will find hard to accept, the converse of this dependence is that each of us can implant homunculi into other people's brains that will enable them to see the world as we see it.

In fact, I can summarise this whole commentary by saying that I have tried to implant inference-detecting homunculi into the brains of those who think about filling-in, and have also tried to implant the homunculus who looks for the origin of other homunculi.

NOTE
1. From Barlow (1995, p. 428). The original had *in some single neuron* between *cause* and *would,* but this is an irrelevant distraction here.

Dennett versus Gibson

Alex Byrne

Department of Linguistics and Philosophy, MIT, Cambridge, MA 02139.
abyrne@mit.edu   web.mit.edu/philos/www/byrne.html

Abstract: Pessoa et al. misinterpret some of Dennett's discussion of filling-in. Their argument against the representational conception of vision and for a Gibsonian alternative is also flawed.

Dennett claims, of his Marilyn wallpaper example (see sect. 9.2 of the target article), that the "brain just represents *that* there are hundreds of identical Marilyns" (1992, p. 46), and so there is no neural representation that has hundreds of Marilyn-representations as parts (as it might be, a "high-resolution bit map"). He comments: "Of course, it does not seem that way to you. It seems to you as if you are actually seeing hundreds of identical Marilyns" (p. 46). The point of the latter remark is that there is a persistent tendency to confuse a *representation of detail* with a *detailed representation.* This is a special case of confusing represented properties with properties of representations (Dennett 1991, pp. 147–48). Because the experience of the Marilyn wallpaper represents it as covered with detailed Marilyns, the temptation is to suppose that the representation itself contains detailed representations of Marilyns. The mistake that Dennett is exposing is thus an intellectual error we make when reflecting on our experience; it is *not* any kind of perceptual illusion.

There is another mistake, according to Dennett: "The hundreds of Marilyns in the wallpaper seem to be present in your experience, seem to be in your mind, not just on the wall" (1991, pp. 359–60; quoted in the target article, sect. 9.2, para. 5). Here Dennett's point is that there is a persistent tendency to invoke something like sense data as the evidence that must be presented to an "inner witness" in the brain. This comes out particularly clearly in the section following the quotation, where Dennett's stooge Otto, looking at a color version of Pessoa et al.'s Figure 1, would suppose that there really is a reddish diamond (not on the page, but presumably in his mind). In the Marilyn case, the temptation is to suppose that there are Marilyn sense data. Again, the mistake is an intellectual error, not a perceptual illusion. (Of course, neon color spreading, unlike seeing Marilyn wallpaper, *does* involve a perceptual illusion, but the perceptual illusion is not the mistake Dennett is identifying.)

Hence, Pessoa et al. have misunderstood Dennett when they say that he erroneously concludes, both times, "that the person's experience of the Marilyns is mistaken" (sect. 9.2, para. 5). (Here

Pessoa et al. must mean that Dennett erroneously concludes that the Marilyn case involves a sort of perceptual illusion, for they evidently agree that the alleged intellectual errors really *are* errors.)

Pessoa et al. think Dennett draws the erroneous conclusion – that the Marilyn case involves perceptual illusion – by reasoning *correctly* from a crucial assumption, the "uniformity thesis." That is why they side with Gibson (1979), and against Dennett (1978), in rejecting the uniformity thesis.[1] So, although these exegetical matters are not significant, what is of some importance is Pessoa et al.'s argument that the uniformity thesis is false because it has false consequences.

According to the uniformity thesis, "perceptual content at the personal level just is the content of brain states at the subpersonal level" (sect. 9.2, para. 4). In other words: when a person has a perceptual experience with such-and-such propositional content, this fact is constituted by the presence, in some part of the person, of a certain sort of representation with that content (as it might be, by the tokening of a neural representation in the subject's "vision box").[2] The uniformity thesis is therefore entailed by the representational theory of mind.

Pessoa et al. think the uniformity thesis leads to trouble in the Marilyn case for two reasons. First, they claim that the thesis entails (or at least motivates) the conclusion that "in the absence of a brain-level pictorial representation of each of the identical Marilyns the person cannot have a percept with the content that there are hundreds of identical Marilyns" (sect. 9.2, para. 4).[3]

(Pessoa et al. interpret Dennett as drawing the conclusion just mentioned. But this does not make any sense. For Dennett certainly holds that there are no such pictorial representations, and that the content of the experience *is* that there are hundreds of identical Marilyns. If he went on to draw the above conclusion, his position would be glaringly inconsistent.)

This parenthetical difficulty aside, it is obvious that the uniformity thesis has no such unwanted consequence. According to it, in order to have an experience with the content that there are hundreds of identical Marilyns, one just needs an appropriate inner representation with exactly that content, and plainly that does not require a representation, pictorial or otherwise, of each Marilyn.

Second, Pessoa et al. claim that the uniformity thesis entails (or motivates) the conclusion that, when looking at the Marilyn wallpaper, "it seems to one that the Marilyns are all there in one's mind or brain" (sect. 9.2, para. 5). Since there are no such pictures (at any rate according to Dennett), the person's experience must, again, be illusory. Pessoa et al. say that "the reasoning depends on the idea that visual experience is pictorial, in the sense that to have a visual experience that is really of hundreds of identical Marilyns is to have a picture in the mind or brain with precisely that content," which in turn "depends on" the uniformity thesis (sect. 9.2, para. 5). But the uniformity thesis is clearly not guilty. First, it does not lend support to the view that inner representations are pictorial (as opposed to, say, sentential). Second, even if we grant that the representation is pictorial, the uniformity thesis does not imply that the person is aware of this inner picture, only that the content of the person's experience is the content of the appropriate inner representation. In the example, the content is *that there are hundreds of Marilyns* (more specifically: hundreds of Marilyns *on the wall*). To think that the person must be aware of the representation itself would be to confuse *what is represented* (what the person is aware of: Marilyns on the wall) with *what does the representing* (a neural picture, which the person need not be aware of). It would be to confuse, that is, represented properties (e.g., being a Marilyn on the wall) with properties of representations (e.g., being a neural picture of a Marilyn). As noted, Dennett himself has emphasized the distinction.

Pessoa et al. correctly state that "at the personal level, there are no . . . representations in visual perception; there is simply experience of the world" (sect. 9.2, penultimate para.; see also sect. 9.1 and Thompson 1995, p. 235).[4] What I do not understand is why this shows that Dennett, and most of us, are wrong, and Gibson is right.

The functional effects of modal versus amodal filling-in

Greg Davis and Jon Driver

Department of Psychology, Institute of Cognitive Neuroscience, University College London, London WC1E 6BT, United Kingdom.
gjd1000@cus.cam.ac.uk   j.driver@ucl.ac.uk

Abstract: Comparisons between modally and amodally completed regions show that perceptual filling-in is not merely the ignoring of absences. Illusory filled-in colour arises for modal completion, but not for amodal completion in comparable displays. We find that attention spreads automatically to modally but not amodally completed regions from their inducers, revealing a functional effect of filled-in colour.

The target article provides a useful summary and discussion of the extensive literature on visual "filling-in." Like many others in this field, we were already convinced of the main empirical conclusion, namely, that neural filling-in goes beyond the mere ignoring of absent information, thus challenging Dennett's (1991) account. We were somewhat disappointed that the article focuses mainly on the challenge from *neural* filling-in alone, and less on the equally persuasive challenge from the psychological reality of *perceptual* filling-in itself. This seemed a missed opportunity, particularly given Pessoa et al.'s extensive discussion of how neural activity might best be related to perceptual states (and also given that few researchers would be so interested in neural filling-in if none arose perceptually!). The authors seem to assume that the existence of filling-in at the perceptual level is merely an uncontroversial and theoretically neutral starting point. However, we think that perceptual measures, not only neural measures, can help establish whether (and when) filling-in goes beyond the mere ignoring of absences.

A direct comparison of modal versus amodal completion may be particularly revealing on this issue. As Pessoa et al. note, several recent authors (e.g., Shipley & Kellman 1992) have shown that these two forms of completion can be induced by very similar displays, and also show several intriguing parallels regarding the completed shape that will be perceived. However, a critical difference remains. No illusory colour or brightness is filled in perceptually for amodally completed regions, which are seen as lying behind an occluder. By contrast, modal completion leads to perceptual filling in of colour and brightness for the completed region, seen in the front plane. The accompanying Figure 1 illustrates this, for two displays used in our own work (Davis & Driver 1997; 1998a). Each comprised two grey segments, abutting a white bar on a dark screen (as in A). Only the depth of the white bar differed between the two displays. When the bar appeared to be closer than the grey segments (due to stereoscopic disparity), the latter were amodally completed as a single grey ellipse that was partly occluded by a white bar (much as it appears for the two-

Figure 1 (Davis & Driver). (A) Stimulus pattern employed by Davis and Driver (1997; 1998a) to compare spreading of attention to and from modally – versus amodally – completed regions. When stereoscopic disparity caused the white bar in the stimulus to appear nearer than the grey regions, the grey regions were amodally completed to form a partly occluded ellipse that continued behind the white bar. Conversely, when disparity signalled instead that the grey regions were nearer to the observer than the white bar, the grey regions were modally completed to form a transparent ellipse that continued in front of the white bar. The region where the modally completed ellipse overlapped with the white bar now appeared filled-in illusory colour. (B) Cartoon of how the stimulus in (A) appeared when the grey regions underwent modal completion to form a transparent ellipse.

dimensional illustration in A). No illusory colour was seen in the amodally completed region itself. By contrast, when the grey regions appeared to be closer than the white bar, they now became *modally* completed, as a continuous transparent grey ellipse in front of the white bar. Critically, illusory colour and brightness was now infused into the completed region, which took on the grey of the curved segments (as cartooned in B), even though no grey was physically present in the centre (just as in A). Thus, modal and amodal completion differ in terms of perceptual filling-in, with illusory colour and brightness arising only in the former case. This difference is not readily accounted for by mere "ignoring of absences," given the similar absences for the two cases. As Pessoa et al. discuss, there have been several previous neuroscience studies on amodal completion, and also on modal completion but none has ever compared them directly, even though this would be particularly revealing regarding the neural basis of perceptual filling-in.

In several recent articles, we have sought to compare the functional effects of modal and amodal completion directly (Davis & Driver 1994; 1997; 1998a, 1998b; Mattingley et al. 1997) as Pessoa et al. recommend. Two studies used displays like those in Figure 1, to examine the effects of modal versus amodal completion upon the distribution of visual attention. Davis and Driver (1998a) found that when attention was drawn to the inducing grey segments (by a sudden change in their size) this also drew attention to the completed region between them in the modal situation. However, this did not arise for the comparable amodal case, thus revealing a difference in the functional effects of the two forms of completion. Davis and Driver (1997) used a measure of distractor interference to confirm that this spreading of attention between inducers and completed regions, for the modal case only, happens even when counter to observers' intentions.

We concluded that the filling-in of colour and brightness, at modally completed regions only, was responsible for this spreading of attention. More generally, we proposed that the presence of filled-in colour at modally completed regions signals that they belong to an unoccluded surface (even when physically absent in the stimulation, due to coincidental camouflage or impoverished illumination; see Davis & Driver 1998a). By contrast, the absence of any colour coding at all for amodally completed regions signals that they are occluded from view. The distinction between occluded and nonoccluded regions of objects is crucial for visual ob-

ject recognition (Nakayama et al. 1989). Moreover, it is also crucial for directing attention towards potentially relevant versus entirely irrelevant information when judging a particular object. Any visible features at occluded (and amodally completed) regions of a relevant object cannot reflect properties of that object, but only those of the occluder. This restriction does not apply to visible features that form part of a camouflaged by unoccluded object, as in modally completed regions. Hence it makes excellent functional sense that attention should spread from inducers to modally completed regions, but not from inducers to comparable amodally completed regions, exactly as we find (Davis & Driver 1997; 1998a).

Our work thus identifies a difference in the functional effects of modal versus amodal completion, using perceptual rather than neural measures. Moreover, the observed difference makes good sense in terms of the particular task faced by the visual system when required to attend to potentially relevant information and ignore irrelevant information. In this restricted sense, we are sympathetic to Pessoa et al.'s enthusiasm for "task level" conceptions of vision in terms of the function(s) served. However, we take issue with their more general advocacy of a "personal level" or "animal centered" approach to vision as a panacea for philosophical confusions in this area. Arguing generically that the "animal . . . simply sees aspects of the world" (sect. 9.1, para. 4) does not advance our understanding of how the animal achieves this: it merely replaces the homuncular little-animal-inside-the-head as an all powerful agent with the larger whole animal. Finally, some of the criticisms of the Marr (1982) approach seemed inappropriate. There are emphatically no homunculi in any of Marr's computer vision systems; and his work was characterised by careful analyses of the tasks to be solved by particular visual processes.

No bridge over the stream of consciousness

Daniel C. Dennett

Center for Cognitive Studies, Tufts University, Medford, MA 02155.
ddennett@tufts.edu   www.tufts.edu/as.cogstud/mainpg.htm

Abstract: Pessoa et al.'s target article shows that although filling-in of various kinds does appear to occur in the brain, it is not required in order to furnish a "bridge locus" where neural events are "isomorphic" to the features of visual consciousness. Some recently uncovered completion phenomena may well play a crucial role in the elaboration of normal visual experience, but others occur too slowly to contribute to normal visual content.

I find this a very useful essay, a model, in fact, for philosophers who want to make a substantive contribution to cognitive science: it makes sense of controversy, dispels confusions, and sharpens our understanding of the more distant implications of a wide variety of current empirical work, an important task that is typically beyond the aspirations (if not the talents) of those working in the labs.

First, let me acknowledge that Pessoa et al. have corrected some errors on my part, errors that betrayed my ignorance of a wealth of empirical and theoretical work that had already been undertaken on the vexed issue of filling-in. I am glad they also pointed out, however, that I did point to the very sorts of empirical experiments that "would" disprove my hunch – some of which should already have been known to me. I was not entirely the innovator I took myself to be, then, but I will settle for the role of catalyst, since my rash interloping has served to direct attention and begin the task of clarification that Pessoa et al. are continuing.

Second, I want to propose a friendly amendment to Pessoa et al.'s fine discussion of what they take to be the fundamental mistake of "analytical isomorphism": supposing there must be a "bridge locus" where features of experience and features of neural

activity are "isomorphic." They are right: the widely felt necessity for just such a bridge locus is a crippling theorist's illusion. But they do not go on to address the natural question of what should replace this overwrought idea. One is inclined to think that something has to be "isomorphic" to something happening in the brain – a "neural-perceptual parallelism" of some sort is required – on pain of lapsing into dualism or sheer mystery. What is the minimal form?

According to Müller's (1896) second axiom (a methodological principle, as they say), "perceptual equalities, similarities and differences correspond to neural equalities, similarities, and differences." The saving move, I submit, is to notice that the data from "introspection" that we need to line up with neural data are not necessarily Müller's "perceptual equalities, similarities and differences" but rather subjects' *beliefs* about perceptual equalities, similarities, and differences. This is another way of making Pessoa et al.'s point about the personal-level/subpersonal-level (first developed in Dennett 1969). What needs accounting for is not necessarily that there is filling-in but that there *seems* to be filling-in – subjects (naively) believe there is filling-in. And explanations in neural terms must be found for the robust provocation of that belief in subjects under various conditions. If we take the manifolds of subjects' beliefs about their conscious states (what I call the subjects' *heterophenomenological worlds*) as our data set, then so long as every such belief (or better, difference in belief) is accountable in neural terms, we "save the data" because those beliefs – rather than their putative objects: actual equalities, similarities and differences – are the data. The possible wild heterogeneity of neural conditions for provoking such beliefs is just what excuses us from hunting for a bridge locus. Sometimes, filling-in-beliefs may be the normal outcome of processes which include genuine filling-in, and sometimes not. In either case, we would have a suitable neuroscientific explanation of the experience of filling-in (in the neutral sense that does not presuppose that it is a veridical experience of filling-in), so minimal materialism would be safe, the burden of proof discharged.

Third, although I have a variety of reservations about Pessoa et al.'s discussions of various candidates for "completion phenomena," I will restrict myself to one point. They note that several of the phenomena they discuss (particularly Ramachandran's [see Ramachandran & Gregory 1991] "artificial scotomata") emerge on slow time scales, ranging from hundreds of milliseconds to several seconds; but Pessoa et al. do not make the further point that any effect taking longer than about 200 msec of fixation to develop cannot itself be part of normal vision, since the eye is seldom fixated long enough for such effects to develop. Far from showing that these varieties of supposed filling-in occur in normal vision, these long time courses show that they cannot occur in normal vision – so any sense we have in normal vision that the background, for instance, is filled-in cannot be due to such effects. Of course, such special effects might nevertheless tell us something about what happens during the first hundred milliseconds of a normal fixation, but any such inferences must be treated with great caution.

Finally, I want to say that Pessoa et al.'s discussion of my "more Marilyns" example is more in agreement with my own view than they have realized. I was using the example to argue (by reductio ad absurdum) against the isomorphism principle that it seems to invoke so readily. I was certainly not endorsing it. I say there is something illusory about the ordinary perceptual experience, not because I endorse the isomorphism principle, but because ordinary people do tacitly endorse it, and it leads them to error. I entirely agree that in one sense, there is nothing illusory about the visual experience in this case: the room is papered with identical Marilyns, and that's just what our visual experience is of: a room papered with identical Marilyns. But people are (mis-)inclined to think that this in turn means that representations of all those identical Marilyns must "therefore" be in their brains – it "follows" from their tacitly held isomorphism postulate. My point is that they are very surprised – to the point of incredulity – to learn otherwise. It is, if you like, a theorist's illusion, but it turns out that we

are all theorists. That is, we tend to assume the isomorphism principle tacitly, and hence are driven to expect that there is more in the brain than there has to be.

Linking spread of neural activity and filling-in: A few more arguments in favor

Peter De Weerd

Laboratory of Brain and Cognition, National Institute of Mental Health, Bethesda, MD 20892. **pdw@ln.nimh.nih.gov**

Abstract: This commentary sides with Pessoa and his colleagues in arguing that some types of perceptual filling-in are linked with a spread of cortical activity, a hypothesis that has often been rejected on philosophical grounds. Some recent data are discussed that strengthen this linking hypothesis and indicate that a spread of cortical activity may be essential for normal surface perception.

While reading this commentary, you probably perceive an entire page filled with English words. The homogeneity of this perception is at odds with the fact that, because of acuity limitations, you would not have noticed whether this paragraph had ended in another language. Our perception of objects (in this case a page of text) is thus determined not only by visual processing and its limitations, but also by expectations built through prior interactions with similar objects. Perceptions of the world are not a reflection of some "objective" reality, and distributed patterns of brain activity correlated with our perceptions do not simply represent such a reality. In this context, Pessoa et al. provide a thought-provoking discussion of the validity of linking propositions used to explain several visual illusions.

Pessoa's discussion of linking hypotheses is especially pertinent to perceptual filling-in, which can be experienced, for example, by maintained peripheral viewing of a gray square surrounded by dynamic texture. Dennett (1991) suggests that the resulting filling-in of the gray region by the texture reflects a symbolic operation by which texture is assigned to the gray region, rather than an isomorphic spread across visual cortex of activity representing the texture. The latter hypothesis is considered untenable because it suggests that perception is caused by a homunculus viewing a final, image-like representation of the "objective" world (sect. 5). Pessoa argues that it is an empirical matter whether cortical spread correlates with perceptual filling-in, and that the possibility of isomorphism should be disconnected from the idea of a final representation (sect. 8). In this respect, it is useful to consider that the visual system is hierarchically organized. Lower-order areas such as V1, V2, and V3, where elementary features are analyzed, such as brightness and texture, are retinotopically organized; neighboring points in the visual field are represented by neighboring points on the cortex. In addition, many cortical neurons show lateral connections extending for several millimeters across the cortical mantle. Thus, the anatomy of lower-order areas makes lateral neural spread a likely mechanism for the perceptual spread of features such as brightness and texture. It is unknown how perceptual filling-in affects activity in higher-order areas that are not retinotopically organized, but any neural correlate of featural filling-in in those areas is unlikely to be isomorphic with the percept. In fact, higher-order areas might be involved in more symbolic types of filling-in, similar to the operation by which you perceived the entire text as consisting of English words when you started reading. Thus, percepts are correlated with activity in many cortical areas simultaneously, and whether there is isomorphism between the percept and neural activity is a by-product of the anatomical organization of those areas, and of the domain in which the filling-in occurs.

With this in mind, Pessoa and his colleagues encourage the exploration of specific linking hypotheses (sect. 10), but also point to potential pitfalls (sect. 4.5). Neurophysiological studies often com-

pare single-neuron recordings in animals with human perception. This is a big step to take, especially if the animal is anesthetized and not perceiving anything. We have used anesthetized Cebus monkeys (unpublished experiments conducted with Ricardo Gattass at the Federal University of Rio de Janeiro, Brazil) in an attempt to replicate the increased ("climbing") activity in V2 and V3 neurons of awake Rhesus monkeys that correlated with perceptual filling-in (De Weerd et al. 1995). We did not find any evidence of climbing activity in V2 and V3 of anesthetized monkeys under conditions that caused perceptual filling-in in human observers, and climbing activity in neurons of awake Rhesus monkeys. This negative result might be specific to the Cebus monkey, or to the anesthetic agents used, but the most parsimonious conclusion is that human conscious perception often cannot be compared with neural activity in unconscious monkeys. The absence of climbing activity during unconsciousness enhances its possible role as a neural substrate for perceptual filling-in.

Pessoa and his colleagues did not address why different types of filling-in follow different time courses (sect. 4.5) and thereby miss the opportunity to explain the relevance of filling in illusions for studying filling-in processes during normal surface perception.[1] Why would brightness spread within milliseconds (Paradiso & Nakayama 1991) and texture only after seconds (De Weerd et al. 1995)? This is because in the former study, the inward spread of brightness was measured within figures with existing boundaries, whereas in the latter, boundary representation separating gray and textured regions had to adapt before filling-in could occur. Recent psychophysical data from De Weerd et al. (1998) support that the time required for the filling-in of a gray region by a surrounding texture did not reflect a slow filling-in process, but rather a slow adaptation of boundaries followed by a fast featural spread. Hence, the climbing activity in V2 and V3 during perceptual filling-in may reveal mechanisms involved in fast filling-in processes during normal surface perception. To further investigate this linking hypothesis, measures of perceptual filling in must be correlated with fast changes in neural activity on a trial-by-trial basis within the same monkey.

NOTE

1. If perception by definition never reflects an "objective" reality, then the word "illusion" is never appropriate. However, the perceptual filling in of a gray region by surrounding texture during a maintained peripheral viewing reflects properties of the visual system, rather than a physical change of the stimulus; in this limited context it can rightly be referred to as a visual illusion.

Area, surface, and contour: Psychophysical correlates of three classes of pictorial completion

Birgitta Dresp

Laboratoire de Psychophysique Sensorielle, EP 618 du C.N.R.S., Université Louis Pasteur, 67000 Strasbourg, France. **bee@currifl.u-strasbg.fr**

Abstract: A simple working taxonomy with three classes of pictorial completion is proposed as an alternative to Pessoa et al.'s classification: area, surface, and contour completion. The classification is based on psychophysical evidence, not on the different phenomenal attributes of the stimuli, showing that pictorial completion is likely to involve mechanistic interactions in the visual system at different levels of processing. Whether the concept of "filling-in" is an appropriate metaphor for the visual mechanisms that may underlie perceptual completion is questioned.

Pessoa et al.'s target article on picture completion argues in favor of a scientific classification, or "working taxonomy" of the perceptual phenomena, encourages their experimental investigation, and advocates seeking explanations in terms of goal-specific visual processes and, as far as possible, neurophysiological mechanisms.

Holistic theories of perceptual completion such as Dennett's (1992) are reviewed, discussed, and finally rejected.

As a psychophysicist actively involved in the study of completion phenomena, I agree with Pessoa et al.'s conclusion that picture completion involves active visual processes (see also Dresp 1997) for which we can find mechanistic explanations that can eventually be linked to neurophysiological substrates. However, their classification of the various completion phenomena strikes me as unnecessarily complicated. The authors call theirs a "working taxonomy . . . meant as a step toward conceptual and terminological clarification." Two general classes of phenomena are introduced: (1) "amodal completion" versus "modal completion" and (2) "boundary completion" versus "featural completion." Pessoa et al. suggest that these divisions cross-classify each other with no hierarchical organization implied. Examples of each of the four subclasses are then discussed.

I would work with a much simpler classification based on psychophysical evidence rather than phenomenology. Classical as well as recent data on perceptual completion suggest that there are only three types of pictorial completion, each involving a different level of visual or perceptual processing. The evidence also suggests that the distinction between "modal" or "amodal" completion is not very useful. The three "operational" kinds of completion are area, surface, and contour completion.

Area completion. In area completion, simultaneous contrast (Diamond 1953; Dittmers 1920) is diffused by the borders of visual objects and spreads out into areas or regions that have no clearly defined spatial limits or boundaries (Spillmann 1981). It is sometimes accompanied by the emergence of perceived relief and apparent depth as in Mach Bands (Mach 1822). Whether phenomenally emergent or not, area completion can be measured psychophysically by having observers detect a small light target at various spatial positions relative to the object border that induces the spreading out (e.g., Dresp & Bonnet 1993; Fiorentini & Zoli 1966; Wildmann 1974). Detection thresholds gradually decrease with increasing distance of the light target from the object border (e.g., Dresp & Bonnet 1993); at some distance, they drop down to the level of control thresholds, measured on areas that are not "contextually contaminated." These local alterations of visual sensitivity change their sign when the contrast sign in the stimulus changes and are reliable psychophysical correlates of the mechanisms underlying area completion. Lateral interactions between contrast detectors in the retina and the visual cortex have been suggested as a neurophysiological explanation (e.g., Fiorentini 1972; Reid & Shapley 1988).

Surface completion. In surface completion, fragments of real or apparent contours give rise to perceptual closure and make the completed regions emerge as figures from the ground (e.g., Koffka 1935; Metzger 1935). Surface completion may or may not be accompanied by pictorial attributes such as illusory contours, apparent depth, or a particularly strong brightness enhancement on the surface, as in the case of the Kanizsa figures (as in Figure 6 of the target article). The important point is that perceptually closed surfaces, irrespective of whether they are filled by apparent contrast, show phenomenal properties of figural relief or depth, or have illusory contours, produce the same psychophysical effects in a variety of visual tasks. All perceptually closed surfaces can produce apparent motion as well as kinetic and stereokinetic effects (e.g., Mather 1988; Vallortigara 1987; von Grünau 1979): They have shape invariance and view stability, which means that the percept resists changes in spatial orientation and position relative to that of the observer (e.g., Carman & Welch 1992; Nakayama & Shimojo 1992; Vallortigara 1987) and they do not affect visual sensitivity to small contrast signals when these are presented on the surface at a position not too close to its border (Cornsweet & Teller 1965; Van Esen & Novak 1974; Dresp 1992). Although surface completion may in some cases be initiated by the same mechanisms that initiate area completion, the psychophysical data tell us that in the visual system surfaces are likely to be represented at a higher level than areas.

Contour completion. In contour completion, collinear lines, edges, or fragments are perceptually joined or grouped together by the visual system. Collinear stimulus fragments do not have to define a closed surface, or partially delimit an area, to trigger completion. The phenomenal effect of contour completion varies. In some cases, an illusory line may emerge perceptually, as between two collinear edges in half a Kanizsa square; in other cases, not much is visible to the naked eye. Psychophysically, the "joining together" operation that underlies contour completion can be measured between lines and edges of any kind provided they are collinear, which means that they can be perceptually aligned. These results tell us that the thresholds for the detection of small contrast targets are considerably lower when the target is inserted between collinear stimuli (lines or edges) than at any other location in the stimulus display (Dresp 1991; Dresp & Bonnet 1991; 1993; 1995; Dresp & Grossberg 1997; Kapadia et al. 1995; Wehrhahn & Dresp 1998; Yu & Levi 1997). These data show unambiguously that the visual system "expects" something to appear within gaps between collinear fragments and is ready to fill in the missing information. This readiness to respond is reflected by the significantly increased visual sensitivity measured in the psychophysical detection task. Long-range interactions between orientation selective neurons in the visual cortex (e.g., Gilbert & Wiesel 1990) provide a plausible neurophysiological explanation of contour completion across spatial gaps.

Conclusion. The three kinds of completion phenomena described above can help explain pictorial completion. Some of the behavioral consequences of the underlying mechanisms can be measured in psychophysical experiments (i.e., "at the level of what the animal accomplishes in its environment"). The psychophysical reality rather than the phenomenal appearance of visual stimuli should govern our scientific working taxonomy. What should be regarded as a mechanism and what as a phenomenal attribute of the stimulus must be clearly distinguished. If "filling in" is to describe one (or several?) mechanism(s), then the term is tendentious, implying a ready conclusion about the nature of the operation(s) that are used by the visual system to achieve completion. Could the whole debate about filling-in be just another example of a well-discussed, ill-posed question?

In summary, I generally believe that down-to-earth and strictly mechanistic conceptualizations rather than overly sophisticated, phenomenal descriptions make it easier to approach visual percepts from the bottom of their genesis. This might save scientific minds from jumping to conclusions regarding whether or not the brain is jumping to conclusions.

Quasi-modal encounters of the third kind: The filling-in of visual detail

Frank H. Durgin

Department of Psychology, Swarthmore College, Swarthmore, PA 19081.
fdurgin1@swarthmore.edu   www.swarthmore.edu/SocSci/fdurgin1/

Abstract: Although Pessoa et al. imply that many aspects of the filling-in debate may be displaced by a regard for active vision, they remain loyal to naive neural reductionist explanations of certain pieces of psychophysical evidence. Alternative interpretations are provided for two specific examples and a new category of filling-in (of visual detail) is proposed.

> No worse evil can befall a man than to hate argument.
>
> Plato, *Phaedo*

Kanizsa (1979) described amodal perception as encountered, rather than seen. When no-nonsense neuroscientists suggest that the filling-in debate is merely semantic, they miss the point that


Figure 1 (Durgin).   A density aftereffect (see Durgin 1995) can be experienced by first gazing at the top fixation bar for about 30 seconds. Upon glancing at the lower fixation square the lower left texture will appear much less dense than the right.

all language is merely semantic. The interpretation of evidence in the understanding of conscious perceptual experience sometimes requires subtle thought, not simply better recording techniques. The target article has taken a somewhat schizophrenic stance with regard to the consideration of evidence. Pessoa et al. argue in the end that concepts of active vision can replace the debate about filling-in with more worthy questions, yet the main treatment of the psychophysical "evidence" is so reductionist in spirit that I feel compelled to try to add a few twists to what I regard as false dichotomies.

There is more to life than boundaries and surfaces. Although Pessoa et al. decry the legacy of Marr (1982), they are content to retain division of visual representation into boundaries and surfaces. In surface perception they lump brightness and texture and any other "features" of surfaces. Is anything left out? It is my contention that the insight suggested by Gibson (1979) and fleshed out in modern debates about consciousness, is that conscious perceptual experience (despite being mediated by neural processes) is of the world, not of the cortical codes. Gibson's ideas about direct perception seem most applicable to the understanding of perception-action coupling (e.g., Milner & Goodale 1995), but part of the reason the concept of direct perception appeals to college sophomores is that perception sure *seems* direct, to paraphrase Dennett (1991).

There is a third kind of filling-in related to this illusory directness, beyond surfaces and boundaries. It is an aspect of what may be called the Grand Illusion of perception – that the perceived world is complete and fully detailed in all directions, despite the unarguable evidence that the internal representation of the world that we can study in visual cortex is nonuniformly sampled (i.e., blurry and poorly localized in the periphery). This third kind of filling-in, which is largely (and forgivably) ignored by those interested in cortical representation rather than perceptual experience, is the filling in of visual detail.

In the space I have here I will provide one example of the sense in which perceptual experience fills in (and seems to have access to) detail beyond what is actually available to it by any straightfor-

ward neuro-cognitive standard. The example is personally dear to me because its discovery was what first convinced me that perception was interesting. It is a rather strong aftereffect of texture density that can be experienced with the aid of Figure 1. After adapting to a dense texture to one side of fixation, subsequent textures presented in that region appear to have far fewer elements than their identical counterparts presented in an unadapted region. Is this because some of the dots are now missing? Or is it because the information underlying my perception of texture density is actually a scalar value, which is "interpreted" (in Helmholtzian fashion) as indicating the presence of such and such a distribution of dots?

It is convenient to speak of the dots as forming a texture that is represented cortically by means of summary information. But what am I to make of the fact that the individual dots all *seem* (to introspection) to be present in consciousness both before and after adaptation? I suggest that this perceptual encountering of detail is a quasi-modal filling-in that is ubiquitous in the efficient registration of the visual world.

Is there such a thing as quasi-modal perception? One of the favored arguments for neural "filling-in" is the evidence of "completion neurons" firing in the visual blind spot (Fiorani et al. 1992). What the target article neglected to remind its readers is that Fiorani et al. also reported the same sorts of neurons firing in response to completions behind occluders – that is, in cases of amodal completion (see also Gattass et al. 1992). Clearly these neurons are doing something more complicated than painting lines into conscious experience.

A second kind of misleading argument comes from the studies of Murakami (1995) of motion aftereffects in the blind spot. According to this study, perceived motion in the blind spot (under conditions of perceptual completion) is neurally represented, so that it will produce measurable motion aftereffects in the neurons receiving direct input only from the contralateral eye. This interpretation demands unwarranted faith in a high degree of localization of motion information, however. Rebecca Cole and I (Durgin & Cole 1997) have studied interocular texture density aftereffects in the blind spot and found a rather more complicated set of propositions to be true.

As with the motion aftereffect, we found that adaptation to completed textures (across the blind spot) was similar to adaptation to real peripheral textures, and differed from that to a "doughnut" of texture in the periphery. However, we also investigated localization of aftereffect strength and determined that the doughnut condition suppresses aftereffect strength both within the doughnut hole and across its textured doughnut surface. The failure of texture completed at the blind spot to act like a doughnut of texture might therefore be due to (a) rather large peripheral receptive fields for properties such as texture (and motion), and (b) the absence of a region of inhibition (as in the doughnut hole), rather than the presence of interpolated activity (see Hardage & Tyler 1995). As a final twist, we showed that even when the perceptual filling-in of texture into the blind spot was prevented by binocular rivalry, adaptation strength was essentially unaffected (see also Blake & Fox 1974; White et al. 1978).

Although the substrate of perception is cortical patterns of firing, the implicit goal of this substrate is to be transparent in perceptual experience. The quasi-modal perception of detail need not reflect anything like an isomorphic reconstruction, though it sure seems like I can see all the dots.

On the roles of consciousness and representations in visual science

David C. Earle
Department of Psychology, Washington Singer Laboratories, University of Exeter, Exeter EX4 4QG, United Kingdom. **d.c.earle@exeter.ac.uk**

Abstract: It is argued that there is a role for the representational conception of vision, and that this is compatible with the task-level account advocated by Pessoa et al. However, the role of representations must be understood independently of our conscious experience of vision.

Pessoa, Thompson & Noë have argued against computational or representational accounts of vision, which they characterise as "subpersonal," and in favour of a task-level approach at the personal level, which is more concerned with the interaction of the perceiving organism with the world. I shall argue that the two positions are compatible, and that the arguments advanced by Pessoa et al. are based on a misconception about the role of consciousness in vision, a limited analysis of the concept of a representation, and a confusion over the proper description of neural processes in vision.

Our undeniable conscious experience of the visual world gets in the way of an understanding of vision. Paradoxically, the ecologically valid task-oriented approach advocated by Pessoa et al. reduces the importance of the phenomenological experience of vision. If we should be concerned with how the organism interacts with its environment on the basis of visually presented information, as argued by Pessoa et al., then the conscious experience of the perceptual world, the phenomenology of vision, can be regarded as epiphenomenal.

Consider, for example, the case of machine vision: suppose we have a machine that can determine relative depth from disparity and on the basis of this information controls its own movements or those of another machine in the physical environment. We need not speak of consciousness here, yet we have a perceiving "organism" that interacts with the world, whose vision can be understood at the task-level in the way described by Pessoa et al.

Similar arguments can be applied in the case of insect vision: do we need to propose that a fly is conscious of its visual environment? The case of blindsight is also provocative in this respect. If this argument is accepted, then perceptual content at the personal level is concerned not with conscious visual experience, but with the way in which an organism acts on the environment in response to visual information, whether that action is a motor response, for example, landing on the ceiling in response to the rate of optical expansion, or making a behavioural judgement of equality of lightness of two fields.

Arguments about the bridge-locus, analytic isomorphism, and the Cartesian theatre need no longer concern us, at least not in the way discussed by Pessoa et al. The problem of vision can then be understood as being concerned with the relation between the processing of visual information and action in and on the world and not the relation between visual processing and the mental way station of the phenomenology of visual experience (Skinner 1963). The phenomenology of vision may be important in its own right; of course there is something that Jackson's (1982) "Mary" doesn't know, but it has little to do with current visual science.

In this context the role of representations in vision can be more clearly understood. I agree with Pessoa et al. that the function of vision is not to produce an image, a picture in the head, which is then inspected, but some might disagree. For a contrary view, consider the following: "These patients cannot report the colours of common objects from memory (e.g., the colour of a football, cactus, or German Shepherd's back), a task that most people find requires imaging the object in colour" (Farah 1988, p. 312), presumably a Technicolor production by the Cartesian Theatre Company. A representation may be many things: it can be both a physical instantiation of a concept and a description. The concept of money can be represented by almost anything: coins, cows,

time, and so on; and equally, coins, cows, and time are descriptions of money in the appropriate context. One of the pervasive errors in visual science has been to attempt to discover the format of representations – witness the propositional versus analog debate over mental imagery in the 1970s and later (Pylyshyn 1981).

Marr (1982) was careful to define a representation in terms of the information that it makes explicit, be that point by point intensity values in the Grey Level Image, or intensity discontinuities and their geometrical distribution in the Primal Sketch. No homunculus is required to read these representations; a representation is simply the output of one process and the input to another process. The primitives or constituents of these representations are to be understood as assertions about the world, an edge here, a depth value there.

Marr adopted a serial stage analysis of vision, and proposed that the representations were symbolic. We do not need to accept these two points to appreciate the value of representations in vision. In a task-oriented approach to vision, if, for example, the control of an action depends on visual depth information, then that information must be represented in the system in some way. Such depth information is not given directly; it must be extracted from the stimulus and represented so as to guide and control action. That representation may be a symbolic one, as in the 2.5D sketch, or it may be encoded in the activity of populations of disparity tuned neurons in V1. The point is that to say that a neuron is tuned to disparity is to assign an information processing and representational function to that neuron. I can see no inconsistency between this subpersonal approach to vision and the personal task-level approach advocated by Pessoa et al., provided that the goal of the subpersonal level is not taken to be the production of the substrate of the phenomenology of vision. If these points are accepted, then I agree that filling-in does indeed seem a shadow of its former self, but not because it has been shorn of its connections to the representational conception of vision. Rather, it is because filling-in has been shorn of its connections to the conscious experience or phenomenology of vision.

The level of filling-in and when it is cognitive

Richard L. Gregory

Department of Experimental Psychology, University of Bristol, Bristol BS8 1TN United Kingdom. **richard.gregory@bris.ac.uk**

Abstract: This informative and conceptually stimulating target article is very useful. I merely query whether the term "illusory contours" is appropriate for gap filling; "illusory surfaces" seems better – and "fictional surfaces" better still. These seem to be rule based rather than knowledge based, suggesting indeed the importance of distinguishing rules (analogous to syntax in language) from knowledge (equivalent to semantics) for classifying perceptual phenomena.

I find this an unusually interesting, informative, and conceptually stimulating target article about fundamental issues of perception and problems of interpretation of phenomena in terms of often hypothetical neural processing. I would like to comment on only one small aspect: the "level" of illusory surfaces.

The term "illusory contours" is not appropriate when considering, for example, the notion that gaps are filled by a perceptually created nearer surface. Couldn't the term "illusory surfaces" be used?

Referring to section 6.1, I think it is an open question how far down the system it is appropriate to speak of "cognitive processes": Whenever knowledge is involved, I should think. It is surely important to note that object knowledge does not seem to be used for illusory surfaces: missing noses, for example, are not completed. Completion seems to depend on general rules rather than on specific object knowledge. Rules, however, might be regarded as cognitive in the sense that in language both semantics

and syntax are aspects of cognition. It would be interesting to look carefully at the analogy between how we should describe perceptual phenomena and how we describe language. (Incidentally, I have recently introduced the term "sideways" rules to complement "bottom-up" signals and "top-down" object knowledge [Gregory 1997; 1998].)

In the first edition of *Eye and brain* (Gregory 1966), I considered the illusory Phi movement as arising from tolerance to gaps in the normal motion signalling system. But surely the general idea of tolerance, and of the brain ignoring signals, goes back at least to Helmholtz in the last century. This may well seem to be the most economical strategy; but one does need to consider the more general system of which these processes and phenomena are a small part. Once the sewing machine is invented, it is the easiest way of sewing a dress; but it would be absurd to invent the sewing machine for a single garment – or not to use it when it is available!

Filling-in the forms

Stephen Grossberg

Department of Cognitive and Neural Systems, Boston University, Boston, MA 02215. **steve@cns.bu.edu   cns-web.bu.edu/Profiles/Grossberg**

Abstract: Boundary completion and surface filling-in are computationally complementary processes whose multiple processing stages form processing streams that realize a hierarchical resolution of uncertainty. Such complementarity and uncertainty principles provide a new foundation for philosophical discussions about visual perception, and lead to neural explanations of difficult perceptual data.

Pessoa, Thompson, & Noë have provided a timely, scholarly, and persuasive description of phenomena related to filling-in. Having with my colleagues introduced neural models of these phenomena (e.g., Grossberg 1984; 1987a; Grossberg & Mingolla 1985; Grossberg & Todorović 1988), I would like to discuss some basic concepts that the authors have not mentioned.

Surface filling in compensates for variable illumination; it "discounts the illuminant." Discounting the illuminant attenuates brightness and color signals except near regions of sufficiently rapid surface change (Grossberg & Todorović 1988). Without further processing, visual representations would be colored outline cartoons. Filling-in reconstructs a surface representation in the attenuated regions using the surviving color and brightness signals, which are relatively uncontaminated by illuminant variations. This process illustrates the "hierarchical resolution of uncertainty."

Filling-in behaves like a diffusion. Why is this? Filling-in of small objects is restricted to small regions, whereas filling-in of larger objects encompasses larger regions. A short spatial range, like a diffusion, can realize both properties.

What keeps the diffusion within its object? This is accomplished by perceptual boundaries. Topographically organized output signals from the boundary stream to the surface stream define regions in which filling-in is restricted.

Remarkably, boundary completion and surface filling-in are computationally complementary processes (Grossberg 1987a; 1994). As in the Kanizsa triangle of Figure 6, boundaries form inwardly and are oriented between cooperating pairs of inducers. Boundaries also form between inducers that have opposite contrast polarities with respect to their backgrounds. Through the pooling or signals from opposite contrast polarities, boundaries form around an object even if it lies in front of a textured background whose contrasts with respect to the object reverse across space (Grossberg 1997, Fig. 8). This pooling property renders the boundary system output insensitive to contrast polarity. The boundary system hereby loses its ability to represent visible colors or brightnesses; "all boundaries are invisible."

Visibility is a property of the surface system. Surface filling in

spreads outwardly from individual inducers and is unoriented until it is contained by a boundary. Filling-in can lead to visible surface percepts, and is thus sensitive to contrast polarity.

The above properties of boundaries and surfaces are complementary. I claim that brain principles of "hierarchical resolution of uncertainty" and "complementary interstream interactions" reflect the uncertainty and complementarity principles obeyed by the physical world with which brains interact, and to which they adapt (Grossberg 1998b).

People like Dennett, who do not believe in filling-in, face a serious epistemological crisis. Ratliff and Sirovich (1978), for example, have accepted the classical viewpoint that early filtering of visual signals occurs and distorts input luminances while discounting the illuminant, but they do not accept that filling-in also occurs. These early filters generate cusps of activation or inhibition near positions where input contrast changes quickly. Unfortunately, two positions with the same (different) activity in the distorted patterns often have different (the same) brightness in the corresponding percept (see Grossberg & Todorović 1988). Being able to measure equality and difference are fundamental operations in all scientific endeavors. Without filling-in, one gets the wrong answer or no answer at all.

Pessoa et al. mention seeing and thinking, which I discuss as seeing and recognizing (Grossberg 1987a). "Seeing" happens in regions like cortical area V4 and "recognizing" in areas like inferotemporal cortex. How you can recognize something that you cannot see (amodal perception) follows from the properties: Boundaries are invisible within the boundary system (in interblob cortical areas V1 and V2) and (2) pathways from this boundary system to the recognition system can support recognition whether or not the boundary is modally visible within the surface representations in area V4.

A nonclassical definition of a visual illusion is hereby suggested: a visual illusion is an unfamiliar combination of boundary and surface properties. The difference between illusion and reality can also be described in these terms.

Both amodal and modal surface perception are possible. Grossberg (1997) suggested how amodal surface percepts form in the monocular filling-in domain (FIDO), while modal surface percepts form in the binocular FIDO. The monocular FIDO is where monocular brightness and color signals are captured onto surfaces that represent different depths by interacting with compatible boundaries (possibly in the thin stripes of cortical area V2); the binocular FIDO is where these captured signals are binocularly matched and fill in a visible, or modal, 3-D surface representation (probably in area V4) in which figure-ground separation occurs. Both processes are proposed to be accomplished by similar mechanisms (Grossberg 1994), notably double-opponent cells. Double-opponent cells are classically assumed to subserve color percepts. At the binocular FIDOs they do lead to conscious color percepts, but at the monocular FIDOs they do not!

So far as I can see, texture spreading can be understood as a combination of boundary completion and surface filling-in. There is no need to posit a separate mechanism.

Pessoa et al. mention several facts as if they were unrelated: (1) reciprocal bottom-up and top-down cortical pathways; (2) rapid reorganization of receptive fields in response to scotomas; and (3) perceptual grouping in both cortical areas V1 and V2, albeit on a smaller scale in V1. Grossberg (1998b) suggested how the visual cortex integrates perceptual grouping, top-down attention, and stable development and learning by making ingenious use of its laminar and columnar circuits. Once one reaches this level of analysis, debates about "Cartesian materialism" seem to be irrelevant.

Revising locus of the bridge between neuroscience and perception

L. W. Hahn

Department of Psychology, Pennsylvania State University, University Park, PA 16802. **lxh22@psu.edu   psych.la.psu.edu/lhahn/home.html**

Abstract: This commentary proposes keeping the bridge locus construct with a revised definition which requires the bridge locus to be dynamic, representation-independent and influenced by top-down processes. The denial of the uniformity of content thesis is equivalent to dualism. The active perception perspective is a valuable one.

Pessoa, Thompson & Noë provide an excellent summary of the perceptual completion literature and a provocative review of the philosophical constructs underlying it. The most significant contribution is their attempt to unite ideas from neuroscience, psychology, and philosophy in creating a description of perceptual completion. Their conclusions, however, originate from a philosophical perspective and are unlikely to be accepted by the empirically minded vision community.

A revision of the bridge locus construct. In section 8.1, Pessoa et al. reject the bridge locus construct based on a strict interpretation of Teller and Pugh's (1983) concise (perhaps overly concise) definition. However, the bridge locus construct provides a useful focal point at which neuroscience, psychology, and philosophy can converge. A redescription of the bridge locus may make the construct more powerful, and perhaps more acceptable to philosophers.

Let us begin redefining the bridge locus with the premise that the brain is responsible for the experience of visual perception. Because the whole brain is not devoted to visual perception, we can restrict the relationship between the brain and visual perception by including only those neural structures that can be stimulated by a retinal image. Of these neural structures, there must be a subset that is sufficient for the current visual percept. This subset can be further divided into: (1) neural structures directly responsible for the visual percept – the neural percept level – and (2) neural structures that directly and indirectly provide input to the neural percept level. "Top-down" neural structures can provide input to the neural percept level but are not part of the neural percept level. I propose that the neural percept level is the bridge locus: the critical point bridging in Teller and Pugh's (1983) words, "two logically disjoint universes of discourse" (p. 587), physiology and perception.

The proposed definition deviates from previous definitions of bridge locus in five important ways. First, the definition does not require that the neural representation be strictly isomorphic or nonisomorphic to the percept. For example, the bridge loci for the percepts induced by "o" and "**o**" may be the same, although they are clearly discriminable. Second, the bridge locus at any particular time depends on the goals of the observer. For example, the bridge locus for the color red will be triggered by an observer labeling the ink color of the word "green" printed in red, but the bridge locus for the word "green" will be triggered by an observer reading the word. Third, the bridge locus is not restricted to high neural stages. For example, the bridge locus for the letter "i" may occur earlier in the visual system than the bridge locus for the word "I." Fourth, the bridge locus is influenced by contextual information. For example, a letter bridge locus will be triggered by the "0" image embedded in "IS0M0RPHIC" while a number bridge locus will be triggered by the "0" image embedded in "1101011111." Finally, implicit in these examples is a significant role for visual attention. By determining the focal percept, visual attention or "what we are looking for" determines the prevailing bridge locus. By giving this role to attention, a mechanism is provided for determining (1) the location of the "fame" of consciousness (Dennett 1996) and (2) which of the "multiple stages" of filling-in is the useful stage (Ramachandran 1995). Thus, the bridge locus provides an explicit link between perception and

physiology, but to be useful, it must be dynamic, representation-independent and influenced by top-down processes.

Is "uniformity of content" just materialism? In section 9.2, Pessoa et al. encounter the mind-body problem in the form of the uniformity of content thesis – that perceptual content and neural content can be the same. Their rejection of the thesis can be interpreted in two ways. First, their denial of the uniformity of content suggests an endorsement of dualism. They step back from a strict dualist perspective in accepting "the general thesis that facts about brain-level content determine what the person sees or experiences" but they "deny that the general thesis entails the uniformity thesis."

An alternative interpretation is that Pessoa et al. are stressing that neural content is always expressed in neural impulses which cannot be equated with the perceptual content of percepts. This assertion challenges a straw man that would be accepted by few researchers. Perceptual experience is not composed of neural impulses, or, for that matter, spatial frequency decomposition. Nonetheless, overwhelming empirical evidence has demonstrated that perceptual experience is encoded in neural impulses and a spatial frequency decomposition. In summary, denying the uniformity of content is either dualistic or a straw man and, in both cases, highly questionable.

The active aspect of the subpersonal versus personal debate. In section 9, Pessoa et al. argue that visual science is dominated by "subpersonal" neuroscience and psychophysics and that the "personal" task-level or goal-level perspective is neglected. Although their argument that the subpersonal approach is being overutilized is unconvincing, I do agree that perception should not be isolated, but should be viewed as embedded within an interaction between perception and action. How many times must we rediscover the significance of Held and Hein's (1963) finding that visual perception depends critically on the interaction between perception and action? Researchers of eye movements are perhaps the most familiar with the ebb and flow of the perceived importance of this interaction. As Steinman and Levinson (1990) have noted, "we are inclined to believe, based on the history of this area, that interest will last a decade or two to be followed by a period of neglect" (p. 115).

We can't fill in answers to philosophical questions

Lloyd Kaufman

Department of Psychology, New York University, New York, NY 10003.
lk@xp.psych.nyu.edu

Abstract: The target article discusses the classic blind spot, scotomas, subjective contours, and other so-called filling-in phenomena. Its purpose is to evaluate the idea that some theories of filling-in amount to tacit acceptance of Cartesian materialism and a form of psychophysical isomorphism. Pessoa et al. reject what is termed structural isomorphism as well as Cartesian materialism, but claim that neural processes adduced as underlying filling-in may be acceptable without implying isomorphism. The article supports the idea of perceiving as an active constructive process. However, the various subthemes are not clearly related to each other. Topological psychophysical isomorphism is indeed untenable, but the tacit assumption that filling-in enjoys any kind of unique status with respect to illuminating philosophical questions is doubtful.

People have to discover their own blind spots. Some writers believe that stimulating the retina surrounding this spot results in a spreading of neural activity to fill it in. Similarly, when viewing the Kanizsa triangle, we "see" borders but when we examine them closely, they vanish. Nevertheless, we say that the borders are completed (filled-in). Pessoa et al. provide a taxonomy of these and other "completion" phenomena and correctly decide that many of them differ basically from each other. The authors' pur-

pose, however, is to illuminate the issue of the relation between physiological events and conscious experience.

How successful were they? Pessoa et al. discuss various forms of isomorphism. As suggested in Kaufman 1974, Kohler (1929) formulated the doctrine of psychophysical isomorphism to legitimize inferences about the brain based on experiences. If a one-to-one correspondence exists between experienced events and brain events, then one may gain insight into the brain by observing one's inner states; by the same token, one can infer experiences from observed activity of the brain. Kohler lacked the tools in 1929 to study the activity of the brain. If one could tap into the brain and observe its states and their changes, it would be enormously convenient. In the 1940s, with the help of Richard Held, Kohler managed to detect changes in steady potentials at the human scalp and suggested that underlying electric currents in the brain did in fact constitute those brain events that mapped experience. The classic monograph by Kohler and Wallach (1944) on figural after-effects beautifully illustrates the complementary idea that one may make inferences about the brain from experienced events. The shifting in spatial position of contours after adapting to other contours must be due to self-impeding current flow in the fluids bathing the brain. Right?

Wallach (personal communication) thought that Kohler's field theory was not justified by the data he and Kohler had collected. He was at a loss to explain these strong phenomena, although the theory did a remarkably good job in its day. Now we see that Kohler's field theory is highly improbable. However, it illustrates the idea that it is impossible to uniquely identify brain activity that underlies perceptions by studying perceptions alone. By the same token, it is highly unlikely that one could uniquely determine the character of an ongoing perception of any complexity at all from measures of neural activity.

As I wrote in 1974, and as made amply clear by Pessoa et al., isomorphism is tacitly accepted by many in the field of perception. Pessoa et al. describe the latest wrinkle in the isomorphism story, that of structural isomorphism. Structural isomorphism implies that if one could measure relevant brain activity, a simple transform of the data would yield a picture of the events as perceived by the owner of the brain. This does not seem likely. Nevertheless, many of the experiments described in the target article are apparently designed to determine whether the behavior of neurons is correlated with some aspect of perception, for example, spreading of neural activity. Finding such a correlation does not establish a causal link. Analogy, as intimated in the target article, is not an explanation.

Even as we grant that psychophysical isomorphism in its various forms is inapplicable to perception, we must acknowledge remarkable successes in inferring neural functions from psychophysical evidence. The notion of opponency was formulated in the nineteenth century by Hering, long before anybody had sharpened a microelectrode. Hurvich and Jameson (1957) described the workings of opponent mechanisms in detail based on psychophysical evidence. Similarly, lateral inhibition was postulated long before the work of Hartline (1942) solely on the basis of visually experienced phenomena (Hering 1872/1964). In 1960, von Bekesy conducted a psychophysical experiment on himself and, based on his data, proposed a hypothetical neural unit remarkably similar to the classic "DOG" or "Mexican hat" concentric receptive field which could account both for various skin sensations as well as the so-called Mach bands. His theorizing led to many quantitative theories that clarify why events at one retinal place affect perceptions linked to distant retinal places (Ratliff 1965). When such mechanisms are present there is no need to postulate filling-in as action-at-a-distance between populations of neurons. How do we account for such successes?

I suggest that these successes are related to processes that occur very early in the visual pathway. We can measure the stimuli to decide what they "really are." Then we notice that our perceptions do not appear to match the measured properties of the stimuli. If a stimulus is "filtered" at an early stage, all that the cortex has to work with is the filtered version of the stimulus. In that case

it would be fairly easy to guess at the nature of a transformation that would more or less "work." However, the going gets really rough when an early-occurring linear transformation cannot account for the discrepancy. Complicated nonlinear transformations (as, e.g., in multistability) are difficult to define, and the number of possible solutions to the problem explodes. Add complexities such as effects of instruction, of prior exposure to other stimuli, and of attention and the problem becomes truly enormous. Plausible physiological theories are so difficult to formulate that mentalistic interpretations are bound to arise.

Apparently Pessoa et al. focus on filling-in partly because Dennett (1991) referred to the idea of filling-in as "a dead giveaway of vestigial Cartesian materialism." However, Dennett's assertion, quoted by Pessoa et al., is hardly central to Dennett's main thesis, which concerns the nature of consciousness itself. He acknowledged that filling in is an empirical issue. He is certainly right to question the assumption that filled-in properties of a perception have the same physiological bases as actual properties. It is so easy to see the difference between what is filled in and what is actually present. Consider the blind spot. We must use the visible background information to make some kind of "inference" about the actually invisible portion of the visual field. We discover that the blind spot is there because we can use it to cause objects or parts of objects to literally disappear. At least some of the filling-in phenomena described in the target article do not require that an inner construction take place. Some may or may not be due to a kind of action-at-a-distance. "Finding out" may be as useful a metaphor as "filling-in" or "neural spreading."

A lesion in the visual cortex may result in a blind spot that is not seen. The lesion is an island of death in a living cortical sea. This dead tissue cannot subserve any function. In fact, if it had normally acted to inhibit activity of adjacent cortical regions, those regions become disinhibited. This may be why the receptive fields in those regions seem enlarged. If so, we cannot refer to the measured changes in receptive field dimensions as reflecting something like neural spreading. Rather, excitability of surviving cells is enhanced because of the lack of normally occurring inhibitory processes that serve to modulate overall brain activity. This may have nothing to do with the fact that the blind spots produced by such scotomas are invisible. They are not filled in, but are simply dead. They differ from the naturally occurring blind spot because the latter has no receptors. The former have retinal receptive fields, but the cortical cells that own those receptive fields are absent. Either way, the differences between these blind spots and their surroundings are not directly sensed. We discover them. Koffka (1935) described a shadow of linearly varying darkness. We fail to notice this gradual change in luminance. However, if a stripe divides the shadow into two regions, we perceive a lighter region and a darker region. The abrupt changes in luminance at the edges become detectable. Similarly, the abrupt changes in luminance of the Craik-O'Brien-Cornsweet patterns (Ratliff 1965) are clearly supra threshold, while more gradual changes between them are less likely to be detected. Hence, the changes at the edges are the only available information regarding the brightnesses of the different regions. This too may be likened to a kind of finding-out or, perhaps, a best bet as to what is out there.

I mention these examples merely to emphasize that the target article deals with empirical issues that have no special status with regard to philosophical questions. The dispute with Dennett seems basically trivial. Even Pessoa et al. agree that there is no inner screen with a little green guy looking at it, and that consciousness is not something to be entered. All of these points are well made by Dennett. In 1974 I too urged that perception be viewed as an active process that unfolds over time. There is neither an instant when nor a place where events enter consciousness. Rather, perceptions develop in time. They are constructions made after the fact of stimulation. Consider apparent motion. When an object appears to move from A to C, it is perceived as moving through the spatially intermediate point B prior to reaching C. Yet if C is occluded, the object never leaves A, it simply

flashes on and off. So perceiving the object at B must occur after it is perceived at C! Such simple perceptual events can thus be described as constructed after the facts of stimulation – even the perceived order of those events in time.

I agree with the general idea of perception as an active process, but am at a loss to understand how filling-in exemplifies this point of view. Given enough space I could describe other phenomena better suited to making this point. To conclude, I am somewhat puzzled by this target article. I agree with many of the ideas presented. However, the connections among them are difficult to discern. The portions dealing with filling-in are interesting in themselves. So are the comments on isomorphism. Then we leap to the distinctions among subpersonal and personal and task level perception. However, there is no substantial information on how such distinctions would guide specific research, and, moreover, no links to the material on filling-in.

Surface representation by population coding

Hidehiko Komatsu

Laboratory of Neural Control, National Institute for Physiological Sciences, Myodaiji, Okazaki-shi, Aichi, 444-8585, Japan. **komatsu@nips.ac.jp www.nips.ac.jp/**

Abstract: Although there is empirical evidence of neural filling-in, this does not necessarily entail "isomorphic" theory. Most cortical neurons do not respond to a uniform surface and are instead sensitive to surface size and quality. I propose that a population of such neurons encodes the presence of a surface. This scheme is different from either the "cognitive" or "isomorphic" theories.

As is nicely summarized in Pessoa et al.'s target article, there is empirical evidence supporting the existence of neural filling-in at the retinotopic map of early visual cortical areas during perceptual filling-in (De Weerd et al. 1995; Fiorani et al. 1992; Rossi et al. 1996; see also Komatsu et al. 1996). Does it follow from these studies, however, that the "bridge locus" lies in these early cortical areas? I do not think so. I think the questions addressed in these studies are independent of the question regarding where the "bridge locus" is located. In my view, there are two common beliefs on which these physiological studies are based. The first is that surface perception resulting from perceptual filling-in and ordinary surface perception share common underlying neural mechanisms somewhere in the course of visual processing. Clearly, the mechanisms involved are different at the receptor level, but the question then remains: At which stage in visual processing do they start to share common mechanisms and how does this happen? Ideally, this question might be answered if we recorded and compared the activities of neurons during these two sorts of surface perception at every stage in the processing; however, such an attempt has not yet been made.

The second common belief is that in ordinary surface perception, neuronal activities in the retinotopic map in early cortical areas closely correlate with the percept in the sense of topographic correspondence. For example, when a surface with a certain color and shape appears in the visual field, some classes of neurons located in a region of the retinotopic map topographically corresponding to the surface should be activated. Under these two premises, one feasible approach for physiological experiments to study neural mechanisms of filling-in would be to examine whether or not neural filling-in occurs at the region corresponding to the perceived surface in the retinotopic map of early cortical areas.

Does such empirical evidence of neural filling-in entail "isomorphism"? Pessoa et al. do not think so. They find "analytic isomorphism" problematic because it requires the postulation of a "bridge locus" (sect. 8.1). I agree that evidence of neural filling-in does not entail "isomorphism," but for a different reason. When

one speaks about neural filling-in, it is often assumed that neurons with a small receptive field confined within the surface are activated throughout the region in the retinotopic map corresponding topographically to the perceived surface, producing a roughly continuous propagation of signals. Such a scheme is unlikely to function as a general explanation of what is going on in neural filling-in. First, when we examined whether visual stimuli covering the blind spot (BS) activated the neurons in the retinotopic map of V1 representing the BS, we found that most of the activated neurons had very large receptive fields that covered a large part of the BS and extended beyond it (Komatsu et al. 1995). I think for the case of perceptual filling-in at the BS, neurons with small receptive fields confined within the BS that are activated during perceptual filling-in are rare, if they exist at all. Another important point is that we need to consider carefully what happens in the early cortical areas when a surface is perceived in an ordinary condition. Actually, only a minority of cortical neurons respond to a homogeneous surface (von der Heydt et al. 1996). Many neurons have some suppressive mechanisms within and outside of the receptive field and are sensitive to the size of stimuli. The optimum response occurs to a stimulus with a certain size, which is in most cases smaller than the size of the receptive field (Komatsu & Kinoshita 1997). This fact is neglected by the "isomorphic" theory. How should we take into account this fact to interpret neural filling-in then?

Suppose there is a population of such size-selective neurons, each with a receptive field at the same location in the visual field, but different size preferences. Also, assume that the preferred size of each neuron is known. Then, if the visual responses of the neurons in the population are measured, it should be possible to compute the stimulus size presented at that location. If there is a population of neurons that are sensitive to some surface quality, such as color, luminance, or texture, the activity of a population of such neurons should encode the presence of a surface with that particular quality. Similarly, there are many neurons in extrastriate areas that are sensitive to surface size and quality. According to such a scheme, the population of size-selective neurons at each cortical stage encodes the surface up to the limit of the maximum encodable size specific to each cortical area. An important difference across different areas is that neurons at a higher stage have larger receptive fields and are thus capable of encoding larger surface areas. Such a scheme of surface representation conforms to neither the "cognitive" nor the "isomorphic" theories. It should also be added that, in such a scheme, a trade-off should exist between the spatial resolution and encodable surface size. In order to determine the precise location and size of a given surface, interplay between early cortical areas and higher cortical ones is necessary. This might correspond to the notion in the target article that visual processing has the "highly interactive, context-dependent nature" (sect. 8.1, para. 7).

One final point should be made about the relationship between perceptual filling-in and the retinotopic map in early cortical areas. Two different cases should be distinguished; in one, perceptual filling-in occurs across a region in the visual field where there is no distortion at the corresponding location in the retinotopic map. Filling-in at the BS is an example of this. In another case, reorganization or distortion takes place in the retinotopic map, and this might also result in perceptual filling-in. Receptive field plasticity at a binocularly produced retinal scotoma, as shown by Gilbert and Wiesel (1992) is an example of this second case (see also Ramachandran 1993c). Distinguishing between these two cases is very important conceptually because the underlying mechanisms of perceptual filling-in should be very different. We have recently shown that, for the monocular retinal scotoma, perceptual filling-in occurs without any reorganization of the retinotopic map of V1 (Murakami et al. 1997). Thus, monocular scotoma, unlike binocular scotoma, conforms to the first case.

Blindsight in the blind spot

K. Kranda

Psychophysiological Laboratory, Freie Universität Berlin, D-14057-Berlin, Germany. **ksk@komma.zedat.fu-berlin.de**

Abstract: The filling-in process proposed as a cover up for the existence of the blind spot has some conceptual similarities to blindsight. The perceptual operation of a hypothetical mechanism responsible for filling in represents a logical paradox. The apparent indeterminacy of the percept in the optic-disc region can be tested experimentally by viewing the grating test pattern below.

The conviction that the human spirit can conquer the flesh is past its heyday, even though it occasionally crops up in disguised form. Some *BBS* commentaries on a phenomenon called "blindsight" (Campion et al. 1983) document such tendencies. The apparent ability to "see" without a functioning striate cortex is an exciting option even if the actual percept is rather ambiguous (Weiskrantz 1996). Pessoa et al. review another perceptual phenomenon – the "filling-in" over the retinal blind spot – but without mentioning its conceptual similarity to blindsight.

The essence of all preoccupation with filling-in, even if not explicitly stated, is the feeling that the blind spot is a blemish of the retina. This "accident" of embryogenesis is admittedly not so bad as a damaged striate cortex, but a nuisance nevertheless, so that our brain needs to iron it out. Although many of us may feel instinctively that there is a real need for some "corrective action," there is an apparent lack of consensus about how to describe this process. Apart from the traditional notion of filling-in favoured by Pessoa et al., other ideas suggesting that the brain performs surface interpolation (Durgin et al. 1995) or some kind of "finding out" (Dennett 1992) are in circulation. Our brain supposedly does all that work to fill the optic-disc area with visual information. But the reason we do not notice the blind spot is precisely because that part of the retina conveys no information. Let me advance this argument by suggesting a thought experiment. Imagine what would happen if a neural mechanism responsible for the hypothetical filling-in or finding out process were selectively lesioned. An assumed appearance of a black or grey disc in the visual field represents a logical paradox, because such a percept actually conveys information about a border and contrast where there is none. The damage to this hypothetical mechanism would thus generate information by producing an illusory percept of the blind spot. The opposite assumption, implying that nothing happens, also leads to a paradox, because it makes the very existence of a mechanism responsible for filling-in superfluous. Pessoa et al. also concede that "neural filling-in may not be logically necessary" (sect. 5.2) but, paradoxically, they express the opinion that its existence can be resolved empirically.

The key to solving this apparently intractable problem is to consider how we perceive the world outside our field of vision. We do not seem to miss a hypothetical extra eye in the back of our heads, nor do we extrapolate our proverbial wallpaper Marilyns (Dennett 1991) beyond the normal field of vision. Under normal viewing, the world does not appear as if seen from inside a tunnel or through binoculars where sharp borders separate the perceived from the nonperceived. As our visual acuity progressively deteriorates with retinal eccentricity, the perceived world merges with the nonperceived one. This process is gradual to the point of indeterminacy of what we actually see in this zone of transition. There is no boundary, because to see it we would need photoreceptors on either side of its retinal image, or, in case of a retinal scotoma, a functioning neural structure expecting input from the damaged part of the retina.

The indeterminacy of the percept from the region of the visual field occupied by the blind spot can be demonstrated experimentally. There are many examples supposedly showing that this region assumes the appearance of a light pattern distributed over the surrounding retina, but can the observer be so certain that he can

Figure 1 (Kranda).   Grating pattern demonstrating the indeterminacy of the percept in the blind-spot region.

accurately describe the percept? This can be tested by viewing a pattern (Fig. 1) where two orthogonal square-wave gratings can produce a perceptual rivalry within the blind-spot region of our visual field. What kind of percept would we expect if some filling-in process operated in this case? Would the grating of one orientation dominate or would the two gratings alternate in time or perhaps fuse to form a two-dimensional grid? If we fixate the spot in the middle with the right eye and slowly move our head towards it, then at a distance of about 30 cm the white spot on the right disappears. What we then actually see within that white region is controversial. Most observers, including myself, cannot tell what it is. Our readers may, however, experience other types of percept, of which I would gladly receive reports!

If our brain is not actively filling in or finding out to cope with the unfortunate existence of the blind spot, what does it actually do about it? I do not know; probably nothing at all because, like monkeys but unlike scientists, it has no concept of a blind spot.

Gestalt isomorphism and the primacy of the subjective perceptual experience

Steven Lehar

14 Crocked Lane, Manchester, MA 01944. **slehar@cns.bu.edu**
cns-alumni.bu.edu/pub/slehar/Lehar.html

Abstract: The Gestalt principle of isomorphism reveals the primacy of subjective experience as a valid source of evidence for the information encoded neurophysiologically. This theory invalidates the abstractionist view that the neurophysiological representation can be of lower dimensionality than the percept to which it gives rise.

I see no difference between the original Gestalt formulation of isomorphism, and what Pessoa et al. call "analytic isomorphism" except for the secondary issue that in the former, the "bridge locus" need not imply a physical location anchored to the neural substrate, but can refer to dynamic patterns of energy in that substrate, as in Köhler's Field Theory. The authors are mistaken when

they claim that Köhler's concept of isomorphism did not extend to sensory qualities such as color. Köhler clears up this confusion (Köhler 1947, p. 60) by explaining that by similarity of "structural properties" he means that "physiological events must be variable in just as many directions or 'dimensions' as the [perceived] colors are," and he specifically cites Müller's axiom.

The key insight of Gestalt theory is that when we view the world around us, what we see is not the world itself, but in fact it is *primarily* a percept, that is, an internal data structure active within the brain, and only in secondary fashion is this data structure also similar in certain respects to external objects and surfaces, just as an image on a television screen is first and foremost a pattern of glowing phosphor dots, and only in secondary fashion is it representative of a remote scene. Based on this insight, it is valid to examine the properties of the world we see before us, not as a scientist observing the physical world, but as a *perceptual* scientist observing a rich and complex *internal* representation.

We cannot yet say with certainty exactly how information is encoded in the brain; neither can we say with certainty what the neural correlate of any perceptual experience might be. What we *can* say with certainty, however, according to isomorphism, is that when we see a filled-in percept, whether real or illusory, the *information* apparent in that subjective experience is exactly representative of the information coded in the mechanism of the brain. Hence, whenever we find, as in the many manifestations of filling-in, a disparity between what we perceive and our notions of neural representation, we cannot argue, as Dennett (1992) suggests, that what we perceive is somehow more explicit than the representation on which it is based; instead we must revise our notions of neurophysiological representation to account for the properties of the subjective experience as observed.

The argument that filling-in is not logically necessary, because the world we perceive around us can be encoded in a compressed or abstracted representation, seems to be supported by computer image compression algorithms that eliminate redundancies in images by encoding only higher-order regularities. However, the compressed representation is of lower dimensionality than the spatial world we perceive, in violation of Müller's axiom. Furthermore, as in image processing, a compressed representation is use-

less without a complementary image *decompression* algorithm to unpack the compressed code and express it again in the ecologically useful form of filled-in surfaces and objects. In fact, the many manifestations of filling-in implicate exactly such an explicit completion mechanism because the percept is observed to contain more explicit spatial information than the retinal stimulus on which it is based.

In response to the objection of "Cartesian materialism," there is no need for a miniature observer of the filled-in scene, for that scene is a data structure, like any data in a computer, except that these are expressed in explicit spatial form.

Dennett argues that the neural representation of a filled-in percept need not involve an explicit filled-in representation, but can entail only an *implicit* encoding, such as the edge image of the retinal ganglion cells, from which the filled-in percept could in principle be computed – although this calculation need not, according to Dennett, actually be performed. However, if such an implicit representation were sufficient to account for perception then there would be no need to posit *any* further processing beyond the retinal image, since that image already implicitly encodes everything in the visual scene.

The notion of perception by "ignoring an absence" (Dennett 1992, p. 48) is meaningful only in a system that has already encoded the scene in which the absent feature can be ignored. In the case of the blind spot, that scene is a two-dimensional surface percept, every point of which is experienced individually and simultaneously at distinct locations. The onus is on Pessoa et al. to explain how this two-dimensional structure can be encoded as a nonspatial "absence."

O'Regan (1992) argues that visual saccades sample the external world like a data access of an internal data structure, and that therefore the structural representation of the visual world need not be encoded internally, because that information is immediately available in the world itself. However, the three-dimensional spatial information of the external world is by no means immediately available from glimpses of the world, but requires the most sophisticated and as yet undefined algorithm to read that spatial information from the world. Indeed, isomorphism suggests that the required algorithm involves the construction of a fully spatial three-dimensional internal model of that external world as observed subjectively. The fact that O'Regan's argument seems at all plausible to him is explained by the fact that when O'Regan *thinks* he is observing the external world, he *is* actually accessing an internal spatial model of the world, because the external world itself is beyond direct experience. The reductionist arguments accepted by Pessoa et al. are plausible only in the context of a Naive Realism that confuses the subjectively perceived world with the objective external world.

Representational theory emerges unscathed

Dennis Lomas

Philosophy Program, Department of Theory and Policy Studies, Ontario Institute for Studies in Education, University of Toronto, Toronto, Ontario M5S 1V6, Canada. **dlomas@oise.utoronto.ca**

Abstract: Representationalism emerges unscathed from Pessoa et al.'s attack. They fail to undermine a key reason for its influence: it has the theoretical resources to explain many illusory visual experiences such as illusory contours and features. Furthermore, in attempting to undermine representationalism, the authors seem to erect an unduly inflated distinction between neural accounts of perception and personal-level accounts of perception.

Is Pessoa et al.'s attack on representationalism (sects. 9 and 9.1 in particular, paras. 2–7) well taken? They are attacking representational theory at one of its strengths because it seems able to explain many experiences of illusory contours and features (as well


Figure 1 (Lomas). Coren et al. (1986) have shown that expectations of subjects significantly influence the shape of the illusory contour they experience at the center of this figure.

as many other illusions such as hallucinations). A familiar representational story is that cognition relies on representations "in the head" that contribute to (causally mediate) the perceptual experience of illusions. A couple of examples suffice to make this point. An illusory shape that many people experience at the center of Figure 1 is hard to explain without representational theory. Representational theory seems even more indispensable in accounting for another experience with this figure. Suppose that subjects, before experiencing Figure 1 and based on verbal communication, expect to perceive a particular illusory shape, for example, a square. In this situation, Coren et al. (1986) have shown that expectations of subjects significantly influence the shape of the illusory contour they experience. It is hard to avoid a representational conclusion in assessing both these experiences: An internal representation of the shape is causally mediating the visual experience.

Despite their anti-representationalism, Pessoa et al. do not seem to give any directly anti-representational account of illusory contours and features, the central concern of their paper. They may feel that they respond adequately in section 9.2 to the representationalism. There they discuss the experience of seeing hundreds of Marilyns on a wall, maintaining, against some theorists, that the visual experience is not an illusion. Whatever the merit of their case, representational accounts play a crucial role in cases where illusions undoubtedly are experienced, as in the case of illusory contours and features.

Pessoa et al. attempt to bolster their anti-representationalism by appealing (sect. 9, para. 2) to the distinction between subpersonal accounts, which, they hold, are concerned with neural mechanisms and processes (sect. 9.1, para. 3), and personal-level accounts, which they hold are concerned with the norms of rationality (sect. 9, para. 2). It is not clear, however, that such a distinction usefully applies to the phenomena of filling in. Neural accounts of many of these phenomena surely play a role in accounting for the personal-level experience of the perceiver. At any rate, the authors have presented no alternative explanation of the personal-level experience of perceptual filling in.

Furthermore, it is unclear what role Pessoa et al. want the notion of "norms of rationality" to play in accounts of visual perception. Are they claiming that visual perception is "constrained by norms of rationality" (sect. 9, para. 2)? If so, their position seems to run afoul of many experiences. (If visual perception, in the authors' view, is not governed by the norms of rationality, then what is the relevance of raising this notion?) The Müller-Lyer illusion, which is often cited in this regard, seems to show that visual perception is not governed by the norms of rationality (for example, see Crane 1992, pp. 149 ff.). For, like many illusions – including illusory contours and features – the Müller-Lyer illusion, for the most part, persists despite the belief that it is an illusion.

All in all, Pessoa et al.'s personal-subpersonal distinction seems to create a gulf between neural and personal accounts of perception that belies the role that neural accounts play in explanations of visual experience on the personal level, including the experience of many illusory contours and features.

Pessoa et al.'s homunculus charge against representational accounts (sect. 9.1, para. 4) hits these accounts at one of their weak points. Specifically, the authors level their fire against Marr's (1982) theory. In this theory, because interpretation takes place only after a "bottom-up" construction of a three-dimensional model, the concern is raised that an implied, fully cognitive homunculus is taking on full responsibility for interpretation. (Hence, the problem of infinite regress lurks.) So the authors' assessment of Marr is justified. However, it is not clear that all representational models necessarily have such a problem. If early vision is guided from the outset by conceptual "top-down" processes (a claim that is consistent with, for example, Cavanagh 1991 and 1995), the problem of a fully cognitive homunculus may be avoided.

Pessoa et al.'s attack on Marr motivates their "task-level conception of vision": "the task of vision is not to produce representations from images, but rather to discover through a perceptual system what is present in the world, and where it is" (sect. 9.1, para. 4). It is not clear how this stance bears on representational accounts, because these accounts provide key insights into how the perceptual system allows us to recognize and locate objects. In fact, representational accounts of one kind or another seem to be front runners in the study of object recognition. To cite one example out of many in the vision science literature, Cavanagh (1995) has argued that the recognition of many two-tone pictures can only be explained by appeal to template matching. Representational accounts also pertain to the study of perception of location. Peterson and Gibson (1994), for example, have argued that in some situations shape recognition of contours in early vision contributes to figure-ground separation.

Pessoa et al. seem to have failed to dint the armour of representationalism with regard to the phenomena of filling-in that their paper addresses. Furthermore, the problems the authors raise against representationalism may be remedied without discarding the theory. Finally, their personal–subpersonal distinction – a basis of their criticism of representational theory – seems to inflate unduly the distinction between neural accounts and personal-level accounts of perception.

A retinotopic representation of filling in: Further supporting evidence

Ikuya Murakami

Department of Psychology, Harvard University, Cambridge, MA 02138.
ikuya@wjh.harvard.edu   visionlab.harvard.edu/

Abstract: A few findings from our laboratory are provided as evidence favoring "isomrphism" in filling-in. One is the responsivity of macaque-cortical area V1 cells to a stimulus designed for surface filling-in at the blind spot. Another is a phenomenological observation of motion aftereffect confined within a filled-in surface at the blind spot. Our recent study on the monkey's perception of surface filling-in at a scotoma is also mentioned.

An illusion called X in fact comprises many qualitatively different varieties; to perceive illusion X does not necessarily imply that the same illusion occurs isomorphically in the brain; studies of illusion X in a dark room with a chin rest might miss the real purpose, nature, and mechanism of X in a natural environment. Based on these general considerations, Pessoa et al.'s target article provides us with an extensive and well-organized review of the issue of perceptual filling-in (or completion or whatever), although nothing dazzling really seems to be postulated. I basically agree with the authors' conservative view of the underlying mechanism of filling-in: There is no promising way to test whether there is a pictorial neural representation of filling-in that is directly or indirectly equated with our perception. The essence of this is certainly a broad issue, not specific to filling-in, but to discuss it in a positive and constructive way in relation to filling-in would definitely be helpful.

Indeed, the idea of a direct correspondence – "analytic isomorphism" – is, though conceptually possible, quite confusing if used in the context of "isomorphism versus neglect." Pessoa et al. correctly pointed this out. I would just add that if analytic isomorphism implies a particular locus in the brain, then it lacks physiological support. In contrast with our impression of the world, which has both high resolution and meaningful organization, no visual cortical areas identified so far seem to have both.

The idea of an indirect correspondence or "neural correlate" offers a more practical and testable strategy for solving the "isomorphism versus neglect" problem in certain kinds of filling-in. Although it is still dangerous to call a particular neuron a "filling-in cell" (as Pessoa et al. explicitly indicate), cells showing firing patterns correlated with our perception of filling-in are a likely candidate for its underlying mechanism. After all, retinotopy is roughly preserved at least in primary visual cortex (and beyond, to some extent), and cells in such a visual map that respond only to a particular (opponent pair of) color(s) are considered "color-selective local units," although the physical world is devoid of "color." Why is this valid? Because we know that the spectral detail is lost after metameric preprocessing in the retina. Likewise, if cells behaved as if they could not distinguish a filled-in pattern from a physical pattern, why not infer that some filling-in-like preprocessing has occurred before that stage? I would like to add further supporting evidence for a retinotopic representation of perceptual filling-in.

The first hint came from our physiological study (Komatsu et al. 1996). When stimulated by a uniform white rectangle covering the blind spot while the normal eye was closed, a few cells of the macaque monkey's cortical area in V1 the representation within the blind spot were activated as if there were a visual input inside the blind spot (an input that could never occur). This firing pattern is consistent with what Fiorani et al. (1992) found in *Cebus* monkey area V1 using stimuli for line completion.

The second point is a phenomenological report (Murakami, unpublished observation) obtained during my study on motion aftereffect following adaptation to filled-in motion (Murakami 1995). An adapting stimulus was confined within the blind spot of the right eye, but presented to the left eye. After adaptation, motion aftereffect was tested using a large stimulus covering the blind spot. When only the left (normal) eye was used at the test, a retinotopically confined motion aftereffect occurred in the middle of the test stimulus. When the right eye was used, however, the same percept was obtained: motion aftereffect was confined to the middle of a *filled-in* pattern. To my knowledge, this is the first report that surface filling-in is seen as *nonuniform!* The neglect hypothesis could hardly explain this phenomenon. On the other hand, according to the isomorphism theory, filling-in is retinotopically represented at an early level before motion adaptation alters the activation map.

The final comment does not have much to do with the isomorphism theory, but is aimed to supplement section 4.3 with our recent study of filling-in after monocular focal retinal lesion in the macaque monkey. Cortical reorganization after binocular focal retinal lesion was reported to be a candidate for the underlying mechanism of filling-in at the scotoma (Chino et al. 1992; 1995; Darian-Smith & Gilbert 1995; Gilbert & Wiesel 1992; Kaas et al. 1990). However, the time course of this cortical reorganization could be on the order of months (Darian-Smith & Gilbert 1995; Gilbert & Wiesel 1992). We found that just a few days after a monocular lesion, the monkey perceived filling-in at the corresponding scotoma (Murakami et al. 1997). Furthermore, the time course of filling-in at the scotoma in each trial was very quick: the monkey made a required action instantly (<0.4 sec) after the stimulus onset.

Pessoa et al.'s target article gives us a good guide for studying filling-in, demarcating modal/amodal and contour/surface distinctions, criticizing a misleading confusion of isomorphism with

the idea of a Cartesian theater [See Dennett & Kinsbourne: "Time and the Observer": *BBS* 15(2) 1992.], and emphasizing the importance of interactions between the perceiving animal and its environment. My comments above should be taken as an addendum to the current knowledge of filling-in compiled in the target article, rather than as an attack from a radical isomorphist.

Trading in form for content and taking the sting out of the mind-body problem

Erik Myin

Department of Philosophy (EMEP), Free University Brussels (VUB), Pleinlaan 2, B 1050 Brussels, Belgium. **emyin@vub.ac.be**

Abstract: Analytical isomorphism is an instance of the demand for a transparent relation between vehicle and content, which is central to the mind-body problem. One can abandon transparency without begging the question with regard to the mind-body problem.

Pessoa et al. rightly argue against the thesis of analytical isomorphism, a specific form of the more general idea of what can be called *transparency*, the idea that there has to be some kind of very special, inner relation of similarity between the bearer of a perceptual content – be it an idea, a sense datum or a neural process – and the content itself.

Transparency was an invariant feature of all premodern theories of perception. In the crude atomistic version, the bearer of content was simply a minuscule copy of the perceptual object that entered even more minuscule pores in the perceiver's eye. With Aristotle, the guarantee of transparency became dematerialized: it was the (immaterial) form that became common to perceived object and perceiver (Lindberg 1976). Transparency lost its self-evident status in the seventeenth century, with the emerging distinction between primary and secondary qualities. Qualitative continuity between object and idea was shattered for the secondary qualities, as in the case of color.

Transparency resettled itself immediately, however, now no longer as a reassuring guarantor of the continuity of man and the perceptual world, but as a central manifestation of the most conspicuous symptom of this broken unity, the mind-body problem. Still today, the lack of a transparent relation is involved by those who insist on the existence of an unbridgeable explanatory gap between science and experience, suggesting that it is because of the qualitative difference between bunches of neurons and phenomenal experience that we will never bridge the gap.

This appeal to transparency as a criterion and a necessary condition for consciousness is accepted by those who claim that we are making advances on the mind-body problem because we are discovering more and more neurophysiological structures and processes that are transparently related to phenomenology. Obvious examples are the various claims about the reducibility of color sensations, because of the (alleged) existence of a neurophysiological color space in which phenomenological relations between colors are mirrored by neurophysiological relations (Clark 1993). [See also Thompson et al.: "Ways of Coloring" *BBS* 15(1) 1992 and van Brakel & Saunders: "Are There Nontrivial Constraints on Colour Categorization?" *BBS* 20(2) 1997.] The target article unveils in the filling-in literature another area where the course of empirical theory building is deflected significantly by the attraction of transparency. As Pessoa et al. clearly state, we do not need to reject the *possibility* of transparent relations between neurophysiological bearers of content and phenomenological contents, but its *necessity*. It is possible, and perhaps even probable, that neurophysiological state spaces exist which are structurally isomorphic to color phenomenology. In the same way, whether filling in involves spreading neural activity is an empirical matter. The

fundamental point is that this isomorphism is neither necessary (*pace* the skeptic) nor sufficient (*pace* the reductionist) for getting the represented content into consciousness.

Abandoning transparency takes the sting out of the mind-body problem, although believers in the existence of an explanatory gap will surely point out that it is a case of begging the question. Moreover, it obligates us to search for new ways of relating brain to mind. But if the search for structurally isomorphic items can no longer guide us, what should?

The obvious way to proceed is to trade in form for content: if we want to find the brain basis of consciousness we need not look for structural correlates, but simply for whatever neurophysiological items we think could carry the content that is present in our phenomenological states. How it is made out that a certain structure or process carries a certain content is a matter for empirical science to decide, and neuroscience and psychology have a vast set of techniques, which seem to fit well enough to accomplish the task, as is shown by Pessoa's overview of the recent research on filling-in phenomena.

According to the more plausible accounts of consciousness, even the simplest visual experience carries within it a sense of how its visual content is related to self, both bodily and personal, and time and space, both inner and outer (Damasio 1994; Myin 1998). A natural way to explain this informational richness of personal level experience is to construe it as the result of the holistic interaction of a number of subpersonal contents. None of these contents can be conscious by itself, but together they constitute an experience at the personal level.

Subpersonal representation seems necessary because the represented objects are at a distance. This is true even for the subject's body, which is – from the brain's point of view – too far away to be consulted directly. Emphasizing embodiment in this way, leads, contra the target article, not away from but rather toward representationalism. Given some modularity in the brain or simply the assumption that the representational capacities of confined areas in the brain are limited, this holism of content easily implies physical representations that are distributed, and the rejection of the idea of confined visual areas that form the immediate substrate of consciousness. Now one can meet the challenge of having begged the question by abandoning transparency, for here is an *argument* that shows that even if transparency obtains, it alone cannot form a sufficient condition for consciousness. Even if a local visual representation is isomorphic to the visual content of a certain experience, this alone would not suffice for consciousness, because too little of the actual content of consciousness would be carried by such a solitary representation.

Representations, computation, and inverse ecological optics

Heiko Neumann

Fakultät für Informatik, Abteilung Neuroinformatik, Universität Ulm, D-89069 Ulm, Germany. **hneumann@neuro.informatik.uni-ulm.de**
www.informatik.uni-ulm.de/ni/mitarbeiter/HNeumann.html

Abstract: Implicit and explicit filling-in phenomena should be distinguished. Blind spot phenomena and mechanisms of boundary completion can be accounted for by implicit filling-in. Surface regions are "painted" with perceptual quantities, such as brightness, by explicit filling-in. "Filling-in" and "finding-out" relate to different computational tasks. Mechanisms of purposive computation (e.g., for navigation) evaluate local measurements, thus "finding out"; whereas mechanisms for grasping might require passive reconstruction, thus "filling in."

Introduction. Pessoa, Thompson & Noë provide a comprehensive overview of perceptual completion phenomena. Although I agree with the authors' general claim that there is evidence for filling-in processes in the neural machinery, I suggest some alternative interpretations of specific findings related to filling in. I also discuss the potential capability of such mechanisms in terms of cortical computation.

Cortical representation, computation, and perceptual units. Perceptual completion appears as a result of different processes and stages in neural architecture. Contrary to Pessoa et al.'s proposed taxonomy for several distinct observations, I recommend disentangling various types of filling-in phenomena.

1. There has been no explicit experimental evidence so far that the blind spot is mapped as an explicit "empty region" in V1 that must be filled in in order to generate a continuous representation of surface layout. Instead, data on cortical development and plasticity suggest that the blind spot appears as a singularity in the cortical representation of nested computational maps. Hence, retinal regions surrounding the blind spot are ideally mapped in direct spatial neighborhoods with overlapping receptive fields that virtually bridge this physiological scotoma. As a consequence, a region is implicitly filled with surrounding featural quantities without an active spreading process.

2. Physiological recordings of monkey V2 cell firing demonstrate that these cells are responsive to subjective contours in the stimulus layout. Boundary completion by V2 contour cells is indicated by their selectivity to the presence of inducers at *both* ends of the elongated receptive field. Researchers have suggested that contour completion generates a spatially contiguous representation of contour activity in V1 through an inwardly directed completion between pairs of inducers. Contrary to this, we argue that an oriented long-range integration mechanism in the V1–V2 *forward* pathway is sufficient to explain the data (Neumann et al. 1997). *Feedback* from V2 contour cells to V1 contrast cells helps to selectively enhance those initial V1 responses that are spatially related. Hence convergent afferent projections and activity integration generate an implicitly filled-in contour activity without an explicit spreading mechanism.

3. Mechanisms of surface perception are faced with the problem of generating representations that are invariant under, for example, illumination. Yet no neural theory of surface representation has been developed. It is reasonable to assume that for certain tasks, such as grasping, a contiguous neural representation of surface shape is necessary. This suggests that the assignment of perceptual surface properties requires a mechanism of "filling in" to generate such an explicit layout. So far, demonstrations of an active diffusion-like filling-in mechanism have been restricted to the brightness domain. It is reasonable to assume that comparable mechanisms also exist for other properties, such as color and depth. For stereoscopic depth perception, localized boundaries which provide disparity information are integrated on the basis of nonaccidental properties (Nakayama & Shimojo 1992). The principles of generating coherent depth quantities are consistent with the filling-in interpretation.

Over all, mechanisms of perceptual completion might be divided into *implicit* and *explicit* filling-in phenomena. Afferent projection zones from nonresponsive retinal regions, such as retinal veins and the blind spot might appear as singularities in a non-Cartesian spatially nonhomogeneous representation in visual cortex. Boundary completion might be better characterized by integration and selective enhancement than by directed spreading activity or filling-in. In both cases, a coherent "filled-in" neural representation is generated implicitly by the neural mapping and coupling principles. On the other hand, any two-dimensional entity corresponding to surface regions might be actively filled-in with perceptual quantities such as brightness.

Visual tasks, representations, and inverse ecological optics. Mechanisms of perception are related to different computational tasks. In the case of vision, purposive computations, active selection from the ambient optic array, and passive reconstruction are

sample strategies. A plethora of results suggest that there is no unique type of cortical representation and process. For tasks like visual navigation or obstacle avoidance, "finding-out" by simply selecting results from neural motion estimation is sufficient to generate proper behavioral control signals (compare Srinivasan et al. 1997). On the other hand, for grasping purposes we may need a dense representation of surface properties to guide any fine-tuned manipulation. We have shown (Neumann & Pessoa 1998b) that filling-in, modeled as an active boundary-mediated propagation process, generates such a dense representation of reconstructed surface properties. These findings indicate the need for more modeling in future investigations. I suggest that emphasis should be shifted from the observation that "the job of filling-in is to complete images or representations in the brain" to an investigation of the purpose of computational mechanisms in an ecological sense. Filling-in is a potential candidate to support related computations on parametrized maps in visual and other sensory cortical areas.

No evidence for neural filling-in – vision as an illusion – pinning down "enaction"

J. K. O'Regan

Laboratoire de Psychologie Expérimentale, CNRS, Université René Descartes, EHESS, EPHE 75006 Paris, France. **oregan@ext.jussieu.fr**
pathfinder.cbr.com/people/oregan/oregan.html

Abstract: (1) The purported evidence for neural filling-in is not evidence for filling-in, but just for long-range dynamic interactions. (2) Vision is perhaps not an "illusion," but at any rate it is not "pictorial." (3) The idea of the "world as an outside memory" as well as MacKay's "conditional readiness for action" may help approach an "enactive" theory of vision.

No evidence for neural filling-in. I fear that for many readers, the main conclusion of Pessoa et al.'s timely synthesis will seem to be that there exists good evidence for low-level filling-in mechanisms: after all, neurons exist whose outputs seem to correspond to illusory contours or filled-in scotomas, and psychophysical evidence exists showing that the filling-in process develops over time in the way expected from a completion mechanism.

But this is the wrong conclusion to draw. The evidence does indeed show the existence of these neurons and these dynamic processes, but the trouble is: there is no guarantee that they are actually being used for filling-in!

For illustration, consider a "nonisomorphistic" theory, in which the internal representation which underlies the phenomenology of vision is not a map of the incoming luminance distribution itself, but a map of the *discontinuities* of the luminance distribution, calculated over a range of spatial scales. Suppose the visual stimulus is a horizontal bar crossing the blind spot. At the finest scale, the bar would be represented as the presence of corners and edges. At a coarser scale, each end of the bar would be represented as an "end-of-line." At a still coarser scale, the presence of two aligned end-of-lines would be signaled. The calculation of such a code, involving derivative-like operators over more or less large areas of the visual field, would take time, and develop progressively.

Note that within this theory, *there is no filling-in.* The bar is seen as continuous across the blind spot because to be seen as broken, a discontinuity would have to be signaled, yet the blind spot can code no discontinuity, since it codes nothing at all.

Yet, although there is no filling-in mechanism in this theory, there are nonetheless neurons corresponding to the "completion neurons" found in the literature, which code the presence of two aligned end-of-lines. Furthermore, there is a dynamic process involved in calculating the code, which could quite feasibly produce effects such as those which Pessoa et al. refer to as psychophysical evidence for dynamic filling-in processes.

The point I am making is that the purported evidence for neural

filling-in is actually not evidence for filling-in. This is also noted in Pessoa et al.'s admirable paper, but I fear it will escape most people's notice. If, as the authors say, "filling in seems a shadow of its former self" (sect. 9.1, last para.), why not call it something else, such as "long range dynamic interactions," which does not prejudge the issue of an internal, pictorial representation?

Is vision an "illusion"? Pessoa et al. criticize statements we (O'Regan 1992; O'Regan et al. 1996) and others (Blackmore et al. 1995; Dennett 1991) have made to the effect that "the impression of continuously seeing 'all' of a visual scene may be an illusion." They say that this way of talking seems to assume that visual experience has a pictorial nature, and that it neglects the personal/subpersonal distinction.

I agree that use of the term "illusion" may lend itself to such a misinterpretation. My use of the term is more of a literary device, designed to shock people into realizing that vision is not what they think it is. In fact, I agree entirely with Pessoa et al. that the internal representation of the outside world is probably not of a pictorial nature. This was one of the main points of my "world as an outside memory" theory (O'Regan 1992).

Pinning down "enaction." On the other hand, concerning the importance of the personal/subpersonal distinction, although I am convinced of its relevance, I don't know how this distinction can be knitted into a theory that makes testable predictions. I think a promising related path is to be found in MacKay's (e.g., 1986) idea that vision is "conditional readiness for action," and in the view that the world is like an "external memory" (O'Regan 1992). Under these views, "filling-in" is unnecessary, because the perception of an object *is constituted by* the sensory changes brought about by actions undertaken with respect to that object. For example, seeing the bar in the example above consists in an aggregate of different kinds of knowledge: in addition to the symbolic knowledge of discontinuities similar to that postulated above, there is *knowledge concerning sensorimotor contingencies:* I have the visual sensation of looking at a horizontal bar near the blind spot, when, among other things, I know that if I move my eye leftwards, information from the right end of the line will no longer be registered because it has moved into the blind spot. It is this kind of knowledge that is actually part of seeing the bar and knowing where it is. Were there not this peculiar behavior of the bar near the blind spot, one would be obliged to conclude one was not seeing the bar, but perhaps hallucinating it. In a similar way: if I closed my eyes or moved them and nothing happened to visual sensation, I would be dreaming, not seeing.

An analogy with tactile perception may be helpful. When a blind man feels a cup, he feels a cup despite the fact that there are spaces between his fingers. In fact, he can handle the cup, and put his fingers into positions which straddle the handle, the edge of the cup, and this helps him to recognize the cup. Far from needing to be "filled in," the spaces between his fingers actually help in recognizing the cup.

Similarly, apparent "defects" of the visual apparatus such as the blind spot, nonhomogeneities in the way the retina samples space and color, as well as eye saccades, can actually be used to gain information about an object. They are part of what seeing is. They do not need to be compensated for. The feeling of "seeing" consists in "being at home with" the way sensations change when you "handle" things with your sensory apparatus. These are the ideas I tried to develop in the "world as an outside memory" theory (O'Regan 1992).

The practical and conceptual case against isomorphism: Evolution and homomorphism

Valla Pishva

Center for Cognitive Studies, Tufts University, Medford, MA 02155
vpishva@aol.com   www.tufts.edu/as/cogstud/maingpg.htm

Abstract: The case against analytical isomorphism is made within an evolutionary framework. The relevance to neural filling-in is discussed. Homomorphism is argued for as a conceptually superior substitute for isomorphism, and the implications for the personal/subpersonal distinction are explored.

Pessoa et al.'s rejection of analytical isomorphism and their emphasis on a personal level account of the visual animal represents a conceptual leap forward in bridging the gap between visual science and perceptual processes. In the context of analytical isomorphism, Dennett's (1992) conceptual objection to neural filling-in is well related: once information about the contents of the blind spot has been gleaned from the surrounding visual field, no re-presentation of the information is necessary. In the worst case, such information would be kicked back out for inclusion in the proximal stimulus. Far from disproof through just the homuncular reductio ad absurdum argument, there is no apparent practical utility derived for the system through such re-presentation. Indeed, the case against any such isomorphic "Cartesian theatre" (and for or against different neural bases of filling-in) is strengthened by evolutionary considerations as well as neurophysiological and psychophysical data.

In an evolutionary framework, any neural representation isomorphic to either the proximal stimulus or perceptual content (be it in the form of a "bridge locus" or not) is favored only insofar as it facilitates the further processing of the represented elements. One could further stipulate that the processing/resources required to create such a representation should not exceed its derived utility. Consequently, an isomorphic representation would be more likely to occur earlier in visual processing, as the proximal stimulus provides an obvious template from which a neural representation could be passively "read off." Once higher processing delocalizes such information, putting the puzzle back together requires a significant expenditure, and therefore a compelling incentive. Perception is not such an impetus.

The implication for neural filling-in is that if it is to contribute to any type of isomorphic representation, the expenditure required to obtain the filled-in information must be outweighed by the utility derived from its inclusion in the representation. Consequently, although it is a difficult task to calculate the "derived utility" side of the equation, we can still postulate that neural filling-in to complete an isomorphic representation (whether topographical or otherwise) would be most likely to occur when the complexity of the underlying processes is minimized. In terms of the types of isomorphisms made clear by Gallistel (1990), the isomorphism between the filled-in information and the initial stimulus whether parts of the proximal stimulus or a neural intermediate) would be considered *functional but not computational,* as a matter of evolutionary tendency. Indeed, much of the target article's evidence shows that the basis of neural filling-in is passive, frequently relying on propagation across the area of the visual cortex corresponding to the blind spot that uses little more than the identity operator. "Seeing" Marilyns in my blind spot requires a bit more than this. (Such philosophical reverse-engineering may be distasteful to some, but the practice is commonplace in Darwinian thinking and evolutionary psychology, and all are haunted by Popperian unfalsifiability.)

Part of the conceptual flaw inherent in analytical isomorphism can also be traced to isomorphism alone in visual science. The tie between a representational, subpersonal account of visual science and isomorphism is strong. As the target article shows (sects. 2, 8.1), the utility of isomorphism can only proceed so far in the face of the multiple realizability of perceptual states with respect to

neural states; and an understanding of neural underpinnings that is increasingly delocalized (both temporally and physically). However, the concept of *homomorphism* can do more conceptual work as a substitute for isomorphism.

While isomorphism requires a one to one correspondence between two systems without the loss of information, homomorphism requires only that certain aspects of one system map onto the other. Beyond the banal fact that there is information degradation in a mapping between two nonabstract systems, homomorphisms can represent relations where the noncorresponding elements of each system are themselves separate yet closely related functional aspects of their system. Within visual science this allows for a more malleable description of proximal-stimulus/neural and neural/perceptual relations. For example, homomorphism provides an alternative explanation for the seemingly negative result of Cumming and Friend (1980) referred to in the target article (sect. 7.1). While Pessoa et al. postulate that "the mechanisms involved [in completion] may be at a 'higher' processing level than the ones involved in the effect being probed (the tilt aftereffect)" (sect. 7.1), the mechanisms may be at (an evolutionarily preferable) "lower" level, while the content provided by the mechanisms is homomorphic to the proximal stimulus (and the tilt aftereffect is thus a nonrepresented aspect). Why does the absence of a full isomorphism between neural filling-in and "the real thing" cause Pessoa et al. to postulate a "higher level" of processing to avoid a homomorphic account?

Finally, it is interesting to note that Pessoa et al. quote, and subsequently disagree with, Todorović (1987) as proposing that "conceptually the idea of an isomorphism between certain aspects of neural activity and certain aspects of percepts may be more acceptable [than a nonisomorphic mapping]" (sect. 8.1). A nonisomorphic mapping is indeed not an appealing alternative in cognitive neuroscientific explanations. If there are truly no relations between the neural substrate and either perception or the proximal stimulus, then we are on the verge of invoking the "notorious mind-body problem." But I disagree with Todorović's reduction of what is easily conceived of as a homomorphic mapping to its isomorphic elements. Such a reduction, while perfectly feasible within the subpersonal program, is reflective of the neglect of the personal level to which Pessoa et al. refer.

Homomorphism at the level of visual perception and isomorphism at the level of neural relations/organization are different *in kind.* Holding that they are not would be an instantiation of the uniformity of content thesis.

Visual perception and subjective visual awareness

Antti Revonsuo

Department of Philosophy, Center for Cognitive Neuroscience, University of Turku, FIN-20014 Turku, Finland. **antti.revonsuo@utu.fi**

Abstract: Pessoa et al. fail to make a clear distinction between visual perception and subjective visual awareness. Their most controversial claims, however, concern subjective visual awareness rather than visual perception: visual awareness is externalized to the "personal level," thus denying the view that consciousness is a natural biological phenomenon somehow constructed inside the brain.

In Pessoa et al.'s target article, the concepts of "vision," "visual space," and "visual field" are used without explicitly clarifying whether and in what sense these notions are supposed to refer to subjective visual awareness. A theory of visual perception should

be very clear about the role it gives to subjective consciousness (Smythies 1996). Visual awareness only refers to one specific phenomenon in the perceptual chain; it should be confused neither with the whole chain nor with other parts of it.

Visual awareness can be realized in the absence of perceptual input or motor output. The visual appearance of dreams is for the most part identical with perceptual experience of the waking world (Rechtschaffen & Buchignani 1992). Certain neural mechanisms are necessary for creating visual awareness, others are not. For example, (activation of) the retina and the visual pathway up to the thalamus are neither sufficient nor necessary in producing subjective visual awareness. Adventitiously blind people have visual awareness in their dreams even though their peripheral sensory systems or pathways may have been destroyed decades ago (Kerr 1993). Conversely, when we fall asleep, we become functionally blind even when our eyes are opened: we cannot see bright flashes (Dement 1976) or objects (Rechtschaffen & Foulkes 1965) right before our widely open eyes, even though the retina and the visual pathway at least up to the thalamus still function normally (Wu 1993). When visual awareness is created during REM-sleep, bursts of ponto-geniculo-occipital waves bombard the visual areas of the brain, using the same thalamocortical pathway that perceptual input uses.

In the light of these empirical facts, it is difficult to deny that subjective visual awareness somehow resides in the brain. Pessoa et al. nevertheless seem to deny this. They say that perceptual content only exists at the personal level (the level of organism-environment interaction). If this claim is supposed to be about the phenomenal content of visual awareness, it is obviously wrong. Phenomenal consciousness is brought about by ("subpersonal") causal processes entirely confined inside the brain, as the generation of visual dream phenomenology shows. The point is that while the perceptual chain as a whole is not subpersonal, the brain processes involved in the creation of subjective visual awareness certainly are.

Dennett (1991) insists that the contents of consciousness do not exist in the brain at the subpersonal level (see Revonsuo 1994). However, he is also notorious for denying the existence of subjective phenomenal experience altogether (e.g., Revonsuo 1993; Searle 1997), so his views will not be of much help in clarifying the role of subjective visual awareness in perception. When Dennett talks about contents of consciousness, he means *attributions* of content made from the outside, not subjective phenomenal experience. He denies that brain events could be classified into conscious and nonconscious, thus also rejecting the idea of a bridge locus or stage: for Dennett, consciousness is not a real natural phenomenon inside the brain, but an external attribution made on the basis of objective patterns in organism–environment interaction. The target article seems to reject, along with Dennett, any notion of the borderline inside the brain between nonconscious brain processes and subjective visual phenomenology. If Pessoa et al. really do accept Dennett's view of consciousness, they should openly say so, for then it is futile to expect their theory to acknowledge subjective phenomenal experience at all. If they do not accept it, they should explain how exactly they construe the place of subjective visual awareness in perception, in dreams, and in relation to the brain.

A cognitive neuroscience of consciousness, I propose, should take seriously the following fundamental assumptions: subjective phenomenal consciousness is a real natural phenomenon in the brain. It is not a simple "bridge locus" or "stage," but rather a *level of organization* in a complex biological system. This phenomenal level of organization is realized "off-line" when we dream, but becomes causally modulated by sensory input in waking perception (Revonsuo 1995). It constitutes the brain's phenomenal real-time *model* of the world, which we in everyday thinking naively take as the external world itself, as if physical reality were somehow "directly" perceived. Cognitive neuroscience should try to figure out what the existence of such a phenomenal level of organization in the brain means: At which level of neural organization could it be

realized? Do we have appropriate methods to discover such a level? What would even count as "measuring" or "imaging" it?

Pessoa et al. seem to be unable to make the fundamental distinction between the phenomenal model of the world, constructed by the brain, and the physical world itself. This distinction was described by Köhler as early as 1929 in a paper called "An old pseudoproblem." Pessoa et al. do refer to that paper, but they nevertheless remain in the grips of that pseudoproblem. They recite (sect. 9.2, para. 10-11) that the wallpaper Marilyns do not seem to be present in one's mind, but there on the wall; and that it doesn't seem as if one's brain is doing anything when one perceives. Now the brain's model of the world is not supposed to *seem like* a model or like anything the brain is doing. It is supposed to seem like a *world*. That's how the contents of consciousness are experienced, even during dreaming.

Furthermore, voluntary interaction with the environment is guided by the contents of visual awareness, not directly by the stimulus fields or the external environment. Blindsight patients have lost the ability to interact *voluntarily* with parts of the visual world falling in their blind field; yet visual stimuli unable to modulate visual awareness still guide simple forms of *nonconscious, automatic* action. In REM-sleep behavior disorder (RBD), the mechanisms creating muscle atonia during REM sleep fail, and the patient acts out his dreams (Schenk 1993). Such a patient attempts to interact with *the world currently in phenomenal awareness.* For example, a 73-year-old man with RBD, when dreaming, attempted to catch a running man. His wife reported that he jumped off the end of the bed and awoke on the floor, physically injured (Dyken et al. 1995). If an epidemic of RBD were suddenly to spread, all of us would behave in bizarre ways every night, guided by the contents of our visual awareness.

The world the brain tries to *interact with* and *adapt to* through voluntary behaviour is the world-model in phenomenal consciousness, not the physical world somehow directly acting upon us (Revonsuo 1997). Even a theory primarily interested in organism–environment interaction cannot neglect the critical role that subjective awareness – the brain's real-time model of the world – has in mediating perception of, and interaction with, the physical reality out there.

Filling-in while finding out: Guiding behavior by representing information

William D. Ross

Lincoln Laboratory, MIT, Lexington, MA 02173-9108. **bross@11.mit.edu**

Abstract: Discriminating behavior depends on neural representations in which the sensory activity patterns guiding different responses are decorrelated from one another. Visual information can often be parsimoniously transformed into these behavioral bridge-locus representations within neuro-computational visuo-spatial maps. Isomorphic inverse-optical world representation is not the goal. Nevertheless, such useful transformations can involve neural filling-in. Such a subpersonal representation of information is consistent with personal-level vision theory.

Pessoa et al. organize a thought-provoking review of many data and perspectives on perceptual completion. They conclude that the data support the possibility that neural filling-in serves as an underlying mechanism of perceived filling-in in several visual modalities. However, the authors join Dennett (1992) in cautioning against the fallacies of isomorphistic and representational approaches. Pessoa et al. find that Marr's (1982) world-description goal for vision "seems guilty of the fallacy of supposing that there is a homunculus in the head whose job it is to view the incoming information" (sect. 9.1, para. 4). Instead, they encourage the personal-level perspective that the goal of vision is the parsimonious guidance of behavior. This commentary attempts to recon-

cile an apparent conflict between the authors' empirical and theoretical conclusions by addressing how we can understand the empirical evidence supporting neural filling-in, given that filled in representations are not themselves the goal of visual processing.

Rebuilding a bridge-locus. Natural selection does not directly favor world-representing subpersonal visual processing strategies. Only an individual's behavior, when sufficiently subtle to meet survival and reproductive needs successfully despite the challenges and competition of a particular ecological niche, gives evolutionary advantage. Whether or not perceptual systems solve inverse-optics problems could be irrelevant. Moreover, introspective evidence for perceptual representations can only offer first-order intuitions about mechanisms, not logical constraints on their form.

However, ruling out behaviorally useless representation on an "internal screen" as a goal for vision does not remove the need for those visual distinctions essential for adaptive behavior. The minimum requirement is that the visual system achieve neural representations of incoming visual information in which the activity patterns appropriate for driving different responses are themselves distinguishably decorrelated from one another. Stated simply, different responses to different visual stimuli require discriminating visual representations. To revise Teller and Pugh (1983), *behavioral bridge-loci* must be reached.

The goal is not inverse-optical representations of the world but useful re-presentation of incoming information. Since such representations are, by definition, sufficient to guide behavioral decisions, no homunculus need view them. They meet Barlow's criterion that "there is nothing else 'looking at' or controlling this activity, which must therefore provide a basis for understanding how the brain controls behaviour" (Barlow 1972, p. 380).

Filling in while finding out. The visual system does respond to certain basic visual information in parallel across retinotopic space. Behavioral motivation for this extravagance is not mysterious. Parallel evaluation of signals indicating potential prey, predator, or mate is simply faster than serial. A particular color of a ripe fruit, or the tell-tale orientation of the continuous edge of a predator partially camouflaged in the brush offer urgent behavioral cues.

Owing to the nature of optical projection, detecting such cues is not directly possible by local retinal processing. Such events are signaled by the integration of projected contextual information measurable at neighboring or even disparate regions of the retinal image. Pessoa et al. write: "the task of vision is not to produce representations from images, but rather to discover through a perceptual system what is present in the world and where" (sect. 9.1, para. 4). Neural models which achieve their goal, but by "producing representations" of visual information from retinal images, are described below. These models exhibit neural "filling-in" not as their goal but as a consequence of integrative mechanisms of finding out what is where.

Pessoa et al. argue (sect. 1, para. 5) that for brightness "there is good evidence for neural filling-in that involves spatially propagating activity." Yet, in a theoretical discussion, they quote Ratliff and Sirovich: "It may well be that marked neural activity adjacent to edges [rather than neural filling-in between edges] . . . is, at some level of the visual system, that final stage" (1978, p. 847).

Visual information is carried by illumination intensities varying over many orders of magnitude, demanding that the retina make good use of its limited neural dynamic range through measuring center/surround contrasts. These track the much more limited range of local surface reflectance ratios. As a result, retinal responses only show "marked neural activity adjacent to" surface reflectance edges, but these local contrast measures are not "final-stage" correlates of local surface reflectance cues. Instead, they signal local relations which vary greatly with context.

Discovering information about scenic surface reflectances demands additional processing to integrate the relations coded by purely local contrast measures. Horn's (1974) integration model was perhaps the first neurally plausible model to achieve this goal. Integration works through the propagation of signals communi-

cating visuospatial contrast information, not the gratuitous filling in of already known lightnesses. Nevertheless, integration models, including a recent neural model of structure-sensitive selective contrast integration, can account for filling-in effects as well as other lightness data (Ross & Pessoa 1997).

Detecting the likelihood of oriented scenic structural groupings demands the statistical integration of colinear or curvilinear local image boundary evidence. A cortical interpretation and development of the boundary contour system (BCS) model (Grossberg & Mingolla 1985) shows that such computations of a boundary's *coherent energy* rather than its local energy achieve the goal of detecting scenic structural groupings embedded in incoherent noise (Grossberg et al. 1997). Such integrative structure detection can result in the completion of boundaries (i.e., illusory contours) that are statistically likely even where local boundary energy is completely absent.

This commentary suggests how an ecological perspective on representation offers promise in reconciling the personal/subpersonal inconsistencies in vision science uncovered by the target article.

Filled-in sensations: The primordial species of imagery?

Kevin Sauvé

Department of Physiology and Neuroscience, Center for Neuromagnetism, New York University School of Medicine, New York, NY 10016.
ksauve@cnm1.med.nyu.edu

Abstract: Filled-in sensations exhibit a distinctive *mélange* of causal features, resembling perceptual sensations in some respects and imagery in others. This commentary identifies several of these shared causal features and advances the hypothesis that filled-in sensations may constitute the primordial species of imagery, the evolutionary neurofunctional precursor of paradigmatic forms of imagery.

The distinction between perceptions and imagery is fundamental to any taxonomy of sensations. What causal differences distinguish perceptions from imagery? The following criterion is initially plausible and, I think, widely accepted: Sensory systems begin to generate imagery when they generate sensations that are not directly causally dependent upon sensory receptor activity in the corresponding part of the sensory field.

The visual sensations you experience while opening your eyes and foveating on an object are almost always caused by a direct spatiotemporally contiguous causal path from your retina to visual thalamus and cortex. *Perceptions* can be plausibly construed as requiring just such a direct causal connection to peripheral sensory receptors. In contrast, *imagery,* is characterized as sensations arising largely independently of any direct, immediate causal connection with external stimuli – as during paradigmatic imagery states such as wakeful imaginative reveries and sleep dreams. Thus, an important distinction between perception and imagery can be drawn by considering only the causal history of the sensations in question (independently of any qualitative or functional differences between perceptual and imagistic sensations).

Filled-in sensations create trouble for this apparently clear distinction. For, as Pessoa et al. note, many filling-in processes are *dynamically inaccurate,* as when brightness extends across an expanse of one's visual field (Paradiso & Nakayama 1991) or when static fills over a visual stimulus (Ramachandran & Gregory 1991). Dynamically inaccurate filling-in provokes us conceptually because the filled-in sensations are no longer *directly* caused by external stimuli at the corresponding part of one's retina. In this respect, such filled-in sensations are imagistic.

Despite these similarities, however, filled-in sensations are quite distinct from paradigmatic forms of wakeful imagery and more similar to perceptions in at least three respects:

(1) Filled-in sensations are generally not noticed *as* filled-in; they are typically *used as* accurate representations of the world (unless the subject has antecedent doubts about the accuracy of these sensations).

(2) Filling-in processes are largely automatic and not under the control we exhibit in voluntary wakeful imagery. This automaticity is indicated by several lines of evidence: (a) filling-in can occur with great rapidity (Paradiso & Nakayama 1991); (b) filling-in typically occurs without individuals trying to fill-in, even when complex filled-in sensations require several seconds to be generated (Ramachandran & Gregory 1991); and (c) early visual processing activity, in V1 and V2, has been associated with several filled-in phenomena (Fiorani et al. 1992; Pettet & Gilbert 1992). The target article cites no evidence to suggest that subjects can directly control visual filling-in processes. In being largely beyond a person's control, filled-in sensations are dissimilar from most wakeful imaginative episodes, and are more similar to most perceptions, hallucinations, and dreams.

(3) The qualitative features of filled-in sensations are evidently related intimately to ongoing perceptual sensory activity in immediately surrounding areas of one's visual field (Watanabe & Cavanagh 1991; De Weerd et al. 1995). Thus the qualitative features of filled-in sensations are closely but *indirectly* dependent causally upon ongoing stimuli – a significant difference from paradigmatic forms of imagery.

Thus, filled-in sensations share at least three important features with perceptions: immediate behavioral use, automaticity, and close indirect causal dependence on stimuli. We also noted, however, that the dynamic inaccuracies and relative causal independence of filled-in sensations are features shared more by imagery than perception. Thus, filled-in sensations are evidently an intriguing *mélange* of perception and imagery.

This *mélange* of causal features is entirely reasonable. The major task of sensations is not to represent how the environment was when our peripheral sensory receptors were stimulated but to predict how our environment will be when we move – because behavior designed for our near future will be more successful than behavior designed for our recent past. In addition, systematic neural time constraints *demand* internal prediction: transduction and transmission times mean that sensory inputs will be 15–100 msec old when first arriving in the thalamocortical system. And, on the output side, motor signals travel for 50–100 msec before activating the appropriate muscles. Intervening between input and output, our brains need time – often 300–500 msec – to generate motor responses appropriate to sensory inputs. Thus, our sensorimotor system must predict how the world will be in 400 msec, and not merely how it was 15–100 msec ago.

In general, predictive problems are made tractable by adopting a small set of generally applicable simplifying assumptions. Filling-in processes simplify our sensory representations by increasing the continuity of objects and features over space (e.g., by conjoining lines across one's blind spot or spreading brightness across expanses) and time. During eye-blinks and saccades, for example, visual sensations continue independently of stimuli for periods of over 300 msec. When our sensory systems simplify and predict by filling-in over space and time, we should expect that such sensations will be only indirectly causally dependent on current inputs from peripheral sensory receptors in the relevant portion of one's sensory field.

Our sensory systems are basically predictive systems, and simplification is fundamental to prediction. So there is reason to suspect that simplifying, predictive filling-in processes may have evolved quite early in the history of conscious sensory processes. Moreover, paradigmatic forms of imagery (imaginings and dreams) may have arisen as our filling-in capabilities extended ever further in time, ever more independently of external stimuli. Perhaps some filling-in processes constitute the primordial species of imagery, the neural processes from which wakeful imagination and sleep dreams were spawned.

Spatiotemporal unit formation

Thomas F. Shipley
Department of Psychology, Temple University, Philadelphia, PA 19122.
tshipley@astro.ocis.temple.edu

Abstract: Findings in dynamic unit formation suggest that completion processes reflect the optics of our world. Dynamic unit formation may depend on patterns of motion signals that are consistent with the causes of optical changes. In addition, dynamic completion conforms to a local curvature minimization constraint. Such relational aspects of vision are important to consider in linking perceptual experience and neural activity.

Much of the discussion of filling in has focused on static displays. Perceptual unit formation is also evident in moving displays, and research on surface, boundary, and path completion of moving objects highlights the importance of ecological considerations in understanding the link between neural activity and perceptual experience. At the personal level both modal and amodal completion of objects is observed. A substantial portion of Michotte's original discussion of amodal perception included dynamic displays (Michotte et al. 1964). For example, in "tunneling," an object will be seen as moving along a continuous path behind an occluder when the visible parts of the moving object's path align spatially and temporally. Apparent motion, where an object appears to move back and forth between two locations, is an example of modal completion across time.

It might be tempting to hypothesize that neurons that respond to a moving point (and apparent motion) are the locus for the neural completion of an object's changing location; however, a consideration of the optics of our world suggests that the link between neural and personal is not so simple. Sequential luminance changes on the retina may arise from two different environmental events, the motion of an object or the sequential occlusion of multiple objects or textural elements. In the first case, the local motion signals are related directly to the motion in the world; the motion signals that occur as things appear and disappear, however, are not. Hence the latter motion signals should not serve as representations of object locations in the world. These signals are not random; their spatiotemporal pattern is lawfully related to an aspect of the occluding object – the orientation of the occluding edge (Shipley & Kellman 1994; 1997). Each pair of appearances or disappearances produces a motion signal whose direction and magnitude is a function of the orientation of the occluding edge and the relative position of the changing elements. As a consequence, any pair of motion signals could be used to recover the orientation of the occluding edge. Specifically, if two motion signals, represented by vectors, are positioned so that they have a common origin, the tips of the vectors define the local orientation of the occluding edge. Discriminating between the patterns of motion signals that accompany motion and dynamic occlusion allows the visual system to use the motion signals to identify boundaries of surfaces that are not specified by luminance, texture, or other static differences. Such boundary formation processes could aid in the identification of objects seen while moving through a cluttered world.

The distinction between local element motion and occlusion has been reported by a number of researchers who have found that one of two percepts may be experienced in displays where elements appear and disappear. Depending on stimulus conditions, such as the temporal interval between changes, observers may report either (1) an edge hiding and revealing elements or (2) mo-

tion of individual elements (Shipley & Kellman 1994; Sigman & Rock 1974; Wallach 1935/1996). A parallel dichotomy is observed in static displays where a tangent discontinuity is perceived as either a corner in the world or a consequence of partial occlusion.

There are a number of other parallels between static illusory contours and spatiotemporally defined edges. In both cases the bounded region appears to have a surface quality that differs from the surround (Cicerone et al. 1995; Cunningham et al. 1998), but a surface quality difference is not necessary to see a bounding edge (Kellman & Loukides 1987; Shipley & Kellman 1994). In addition, the filled in edge is almost always seen as smooth. Corners in illusory contours are rarely reported and when shown a spatiotemporally defined hexagon subjects frequently mistake it for a circle (Shipley & Kellman 1994). Minimization of change, both across surfaces and along edges, would appear to be a general property of filling in.

The observation of covariation in perceptual aspects of an event, as well as constraints on how filling in occurs, are relevant to issues of representations in vision. Conceiving of perceptual representations as encoding all aspects of a scene is indeed problematic. A more limited meaning of representation may be appropriate, and may help focus on the importance of relationships. Gallistel (1990) offers a definition of neural representations as events or states of the nervous system that are isomorphic to events or states of the world. This definition does not require all aspects of the world be encoded, only the ones that are critical for action. Unlike many uses of the term representation, this is purely mathematical, so it focuses on what relationships (or functions) present in the represented system are preserved in the representing system (e.g., when numbers are used to represent quantities of objects, some relations such as addition and division are preserved, whereas others, such as whether the objects are soft or hard, are not).

In building a bridge between the personal and neural, it may be useful to try to understand which mathematical relationships apply in the representing system. Considering both how the nervous system may instantiate computations (Gallistel [1990] reviews a number of cases where the nervous system appears to use vector representations and vector addition occurs as a consequence of the spatial structure of the nervous system) and the ecological relationships, which suggest that certain perceptual experiences will be coupled (Epstein [1982] provides a review of cases where two perceptual experiences are clearly linked) should be helpful.

Active vision and the basketball problem

Manish Singh and Donald D. Hoffman
Department of Cognitive Sciences, University of California, Irvine, CA 92697-5100. **msingh@uci.edu   ddhoff@uci.edu**
aris.ss.uci.edu/cogsci/personnel/gstudents/singh/
aris.ss.uci.edu/cogsci/personnel/hoffman/hoffman.html

Abstract: It is fruitful to think of the representational and the organism-centered approaches as complementary levels of analysis, rather than mutually exclusive alternatives. Claims to the contrary by proponents of the organism-centered approach face what we call the "basketball problem."

The target article of Pessoa et al. nicely summarizes a long list of perceptual-completion phenomena. To this list we can add a relatively new phenomenon called "dynamic color spreading" (DCS; Cicerone & Hoffman 1991; Cicerone et al. 1995) in which apparent motion triggers human vision to construct illusory surfaces and contours. Figure 1a shows one frame from a DCS movie, consisting of a square containing randomly placed dots. Most dots are red, but a few (which fall within a virtual disk) are green. Figure

Figure 1 (Singh & Hoffman).   Three frames from a dynamic color-spreading movie.

1b and 1c show two more frames from the same movie. From frame to frame no dots move – but different dots are colored green. Setting these frames in motion triggers the modal perception of a circular and transparent green filter gliding over the red dots. It can also trigger the amodal perception of an opaque green disk gliding behind the red dots. DCS is unique in that the completion of the surface and bounding contour is triggered by apparent motion.

The target article properly warns readers of the dangers of analytic isomorphism and Cartesian materialism, and argues that neural filling in need not imply either. However, its positive contribution to methodology in visual science is less clear. The authors contrast the representational approach to vision with an animal-based or activity-based approach – one in which the focus of study is not internal processing, but the whole animal interacting with its environment – and they suggest that visual science should reject the representational approach in favor of the animal-based approach.

It is not clear why these two approaches must oppose each other. No one denies the importance of understanding vision and cognition within the larger context of an organism interacting with its environment. Indeed, representational questions are often motivated by observations at the organism/personal level: given that an organism can do such and such, what computations and representations might subserve this ability? For example, at the organism level, there is indeed a "perception-action system that enables the animal to visually guide its activity and thereby visually explore its environment." But one needs to ask: What mechanisms make such a perception-action system possible? This is the kind of question that the representational approach addresses. Without answers to such questions, our understanding is bound to remain incomplete. Thus the representational and the organism-centered approaches are more fruitfully thought of as complementary levels of analysis, rather than mutually exclusive alternatives.

Pessoa et al. worry that current visual research focuses on representations and ignores issues at the level of the organism. But this research simply reflects the normal modus operandi of science: It is easier to study small pieces of the puzzle first, and complex interactions later. In this regard, it is surprising that the target article offers no systematic plan for research at the organism level. In what ways would such research differ from current psychophysics? What experiments would we run? What insights might we expect? And how would these insights obviate the need for representational accounts? If the authors wish to advocate an alternative approach for vision research, the burden is on them to detail a concrete plan of research, and show why it might be superior.

This deficit is all the more acute because Pessoa et al. do not make clear whether, by perception at the level of the organism,

they mean (a) perception from a third-person perspective in which we as scientists look at the behavior of an organism and decide that it perceives something, or (b) perception from a first-person "phenomenological" perspective. For example, do they believe that the same methodology would apply in both cases?

Pessoa et al. also claim that the subpersonal account of vision creates conceptual confusions regarding what vision is, and that it is guilty of the homunculus fallacy. Again, no confusion results if one keeps clearly in mind that one is dealing with two different levels of analysis. One may talk either of an organism perceiving and acting in an environment in appropriate ways, or of internal processing that allows for such behavior. With this distinction in mind, it involves no confusion to say that the construction of certain representations is a necessary condition for the organism to see (Hoffman 1998; Singh & Hoffman 1997). Take away those representations and the organism no longer sees. Nor is a homunculus fallacy involved: the representations that are constructed need not be passed on to other "higher centers" for further interpretation – some representations are by themselves sufficient to trigger appropriate responses by the organism. Hence the representational approach to vision in itself entails neither Cartesian materialism, nor analytic isomorphism, nor the homunculus fallacy.

Even if we agree, for the sake of argument, that the need for representations is minimized through reliance on perceptually guided tasks of the animal as a whole, it is nevertheless true that the animal engages in many such tasks, and that it must therefore use many representations. We do not agree, however, that this need is minimized through reliance on perceptually guided tasks. Consider the perceptually guided task of playing professional basketball. One must keep track of one's four teammates, the five opponents, two baskets, the ball, the coach, and a stand full of noisy fans. The burden is on advocates of the activity-based approach to show how this task can be accomplished with minimal representations.

How do we see what is not there?[1]

Lothar Spillmann[a] and John S. Werner[b]

[a]Institute of Biophysics and Radiation Biology, University of Freiburg, D-79104 Freiburg, Germany; [b]Department of Psychology, University of Colorado, Boulder, CO 80309-0345. **spillmann@uni-freiberg.de jwerner@clipr.colorado.edu psych-www.colorado.edu/faculty/werner.html**

Abstract: Pessoa et al. provide a valuable taxonomy of perceptual completion phenomena, but it is not yet clear whether these phenomena are mediated by one kind of neural mechanism or more. We suggest three possible neural mechanisms of long-range interaction to stimulate further

perceptual and neurophysiological investigation of perceptual completion and filling-in.

Why do we perceive extended surfaces, including a uniform visual field, and not just skeletons, as might be expected from our concepts of the classical receptive field? Ganzfeld studies and studies with stabilized images show that we do not maintain brightness and color perception for long. A few minutes after being exposed to a uniform sphere, observers report that a bright field becomes gray and ultimately dark, approaching the *Eigengrau* (Knau & Spillmann 1997). Similarly, with careful fixation, brightness, color, movement, flicker, and even a gray square on a dynamic noise background will become degraded and disappear from view. This fading reflects the fact that the neuronal response dies down quickly when a stimulus remains constant in space and time. With eye movements, spatio-temporal modulation normally occurs at the edge of the stimulus, thereby ensuring sustained and undiminished vision, even of an afterimage. The question is: How is the enclosed area filled-in?

Pessoa et al.'s target article focuses on large scale filling-in of uniform areas and discusses them within the framework of stimuli that cannot be seen (as in the blind spot) and of percepts for which there is no appropriate stimulus (amodal completion). What we see cannot be wrong, even if it is illusory (Metzger 1953). Illusions are undeniable, persist despite superior knowledge, are present in animals, and can even be shown to correlate with responses at the single-cell level. This assumption has been a cornerstone for research on color perception but has never had a comparable impact on the study of form perception. For those of us who have loved visual illusions all along, it is gratifying to note that ultimately these phenomena are considered more than just oddities and that they are increasingly used as noninvasive tools to better understand the workings of the human visual system.

The discussion of linking propositions, bridge loci, neural isomorphism, and so on should not distract from the fact that one can be oblivious to such intricacies and yet strike empirical gold. In the early fifties, when single-cell recording in the visual cortex of the cat began, Richard Jung, Günter Baumgartner, Otto Creutzfeldt, and Otto-Joachim Grüsser in the Freiburg laboratories made a concerted effort to find neuronal "correlates" of perceptual phenomena. Grüsser (1956) defended this approach by pointing out that one need not demonstrate a causal connection between a cell's behavior and the perceptual event; what matters is the spatio-temporal correlation. A large body of publications from that time (mostly in German) strongly suggests that there are many parallels between the firing rate of individual neurons and specific percepts, pointing towards the combined discharge of a larger population of such cells as the putative mechanism. Naively or not, firing patterns were understood as mediating the percept of the stimulus, even if perceiving was actually misperceiving (as in the case of the Hermann grid illusion). Where consciousness arises was not asked.

To give but one example: Baumgartner's finding (1961) that in the cat, on- and off-units produce a spatial response distribution that reflects border contrast, makes it highly plausible that the percepts for "brighter" on the high-luminance edge and "darker" on the low-luminance edge are somehow connected to (and presumably brought about by) different neural activities. Following the early proposals of Mach and Hering in the last century and the epistemological parallelism of the philosopher Nicolai Hartmann, Jung (1973) listed a large variety of neurophysiological-perceptual correlations in his seminal chapter in the *Handbook of Sensory Physiology*. Even the staunchest opponent of neuro-reductionism might find it difficult to attribute these correlations to mere coincidence. The similarities in size between receptive fields of single neurons in monkeys and perceptive fields (as their psychophysical analogs) in human vision imply that many neurons essentially behave like a single one by interacting synergistically to produce a given percept (Spillmann et al. 1987).

The long-range spreading in the perceptual phenomena considered by Pessoa et al. may be analogous to the perception of color on a white field surrounded by a chromatic field, Goethe's colored shadows. Long-range interactions must be assumed to propagate the information from the boundary to the "silent" portions of the stimulus. Considering that the main topic of this target article is filling-in to explain illusory contours (boundary completion) and area enhancement (featural completion), the question might also be asked why under certain conditions a genuine hole, such as the gap in a Landolt C, is not filled-in as it would be if the gap were imaged on a scotoma.

Pessoa et al. remind us that the filling-in of subjective contours may be something different from the filling-in of the physiological blind spot. However, they present no compelling evidence that these diverse phenomena require different neurophysiological mechanisms for their explanation. A central question is whether generalized neural circuits or specialized solutions are needed for the brain to "make something out of nothing." We (Spillmann & Werner 1996) have suggested three candidate mechanisms of long-range neural interaction to account for such diverse phenomena as illusory contours, filling-in of brightness and color in area contrast, and filling-in of the blind spot and scotomata, among others.

One promising candidate to explain filling-in is provided by a feedforward circuit in which the signals from spatially separated receptive fields converge at higher levels. An example is the proposal by Peterhans and von der Heydt (see Baumgartner 1990) that subjective contours result from the convergence of neuronal responses to real contours (e.g., offset grating lines) and a second path that integrates these responses in a direction orthogonal to the stimulus pattern. Signals from the two paths are combined so that the output is indistinguishable from the response to a real line.

A second mechanism for filling-in may be associated with the "gating" of long-range horizontal connections between hypercolumns that are separated by inactive cortical regions. In the absence of primary input (due to retinal lesions and, presumably, uniform retinal stimulation), these horizontal connections may provide links (Wertheimer's *Querfunktionen* 1912) that bridge the gaps between stimulated areas. In this way, the cortical representation of unstimulated regions of visual space can be "assigned" a neuronal state corresponding to the neural activity at the edges.

Finally, global interactions between widely separated areas in the brain may be mediated by the synchronized discharge of neural activity rather than through dedicated circuits. Such binding by re-entrant signals from higher areas may explain some of the Gestalt factors (e.g., common fate) in which stimulus elements spaced across numerous hypercolumns are nevertheless perceived as a whole, by virtue of grouping.

Pessoa et al. have made a convincing case that active neural mechanisms are responsible for a variety of filling-in phenomena. We now look forward to further perceptual, psychophysical, computational, and neurophysiological research to determine where and how the brain accomplishes these tasks.

NOTE
1. Please direct future correspondence to Lothar Spillmann at: Arbeitsgruppe Hirnforschung, Institut für Biophysik, Hansastrasse 9, D-79104 Freiburg, Germany.

In defense of neuro-perceptual isomorphism

Dejan Todorović

Department of Psychology, University of Belgrade, 11000 Belgrade, Serbia, Yugoslavia. **dtodorov@dekart.f.bg.ac.yu**

Abstract: It is argued that the notion of bridge locus is compatible with distributed representation and brain interconnectivity. Isomorphism is not a dogmatic condition on explanatory adequacy but a refutable hypothesis, superior to Dennett's proposed alternatives. The assumption of type-type

neuro-perceptual correspondences is more parsimonious than multiple realizability.

The bridge locus is a hypothesized set of neurons underlying conscious experience. One reason given for its dismissal in the target article is that neural correlates of percepts may be distributed over many brain areas. However, the notion of bridge locus does not preclude this possibility. The existence of topographically organized reciprocal interconnections between visual areas can impose a functional unity upon subsets of neurons belonging to different areas. Two interconnected neurons located in topographically corresponding positions in two anatomically distant cortical areas may be functionally closer than two anatomically nearer neurons in the same area. The aspects of neural activity relevant for percepts may not be based on the "external" geometry of anatomical locations but on the "internal" geometry of functional interconnections. Thus, the interconnectivity of cortical regions, which Pessoa et al. see as another problem for the notion of bridge locus, can in fact be incorporated into it. Nevertheless, some aspects of percepts may be determined predominantly by neurons in a single area.

If bridge locus neurons provide the substrate for an ongoing percept, then the structure of their activity should relate in some way to the structure of the percept. Note that the abstract concept of *structure* can be applied both to percepts and neural activations, thus helping to bridge the mind-brain gap; this is a key idea of isomorphism. For example, in the Craik-O'Brien-Cornsweet effect, the perceptual structure to be accounted for is the perceived *lightness step*: the two rectangles are perceived to have different homogeneous lightness levels. However, available physiological evidence suggests that the spatial activity profile of *retinal* neurons would be roughly in accord with the *cusp-shaped* luminance profile; furthermore, a very similar retinal profile is likely to be induced by *step-shaped* luminance profiles (see Todorović 1987). The problem raised by this neuro-perceptual structural mismatch is best seen by noting that retinal activations of equal strength, corresponding to the interiors of the two rectangles, are associated with different levels of perceived lightness. One way to account for step-shaped percepts induced by luminance cusps and steps is to hypothesize that in both cases the activity profile of the relevant bridge locus neurons is *isomorphic* with the percept, that is, step-shaped. A proposal of how both luminance cusps and steps are first transformed into retinal cusps and subsequently into cortical steps, corresponding to perceived lightness steps, is presented in a model by Grossberg and Todorović (1988), which accounts for a number of lightness perception phenomena. Rossi et al.'s (1996) neurophysiological data on lightness contrast, described in the target article, corroborate the isomorphicist prediction that perceived lightness levels correlate with cortical activity levels.

As against isomorphically guided theorizing, Pessoa et al. claim that Dennett (1992) has uncovered the main conceptual error of this notion: isomorphism, they say, maintains that neural correlates *must be* isomorphic to corresponding percepts, whereas Dennett shows that they *needn't be* by pointing out alternative accounts. This criticism is easily defused by substituting "must be" with "are," thus transforming isomorphism from a purported dogma into an ordinary hypothesis. The proposal that neural correlates *are* isomorphic with percepts, like any other scientific hypothesis, is not endangered simply by indicating that alternative accounts are conceivable, but by demonstrating that they are preferable.

Are they? One of Dennett's alternatives is that the brain ignores the absence of a representation. However, the target article describes several examples which strongly suggest the presence of neural representations in filling-in phenomena, and none that suggest its absence. The other proposed alternative is that the brain "jumps to a conclusion" that a region is colored, or that it "attaches a color label all at once." However, before such proposals are considered as serious candidates for *neural* corre-

lates of percepts, they must be reformulated in neural terms, preferably based on current knowledge of the structure of the visual system and color physiology. For example, what happens to neurons in the portions of visual cortical areas activated by the two rectangles when the brain "arrives at the judgement" that constitutes our seeing of particular lightness levels in the Craik-O'Brien-Cornsweet display? A neuro-perceptual theory requires a neuro-perceptual vocabulary. Isomorphism may be a simplistic, "pictorial," or just a plain wrong answer, but at least it is an answer to the problem of the relation of neural and perceptual states, as it involves refutable hypotheses about neural activity. In contrast, talk of judgments, conclusions, labeling, and the like, without explaining how such processes are neurally implemented, does not constitute an acceptably formulated answer, neither does any account that side-steps these issues, such as Pessoa et al.'s task-level approach. Furthermore, their "fundamental conceptual point," which disallows inferences from perceptual content to neural medium, appears to rule out, purely on "logical, conceptual or methodological grounds" such brilliant speculations as the Young-Helmholtz inference from color-mixing data to cone spectral sensitivities, or Mach's inference from the appearance of his bands to the existence of retinal lateral cross-connections.

Pessoa et al. note that the isomorphicist approach, as outlined by Todorović (1987), would be plausible only in the case of type-type neuro-perceptual identities, but not if the same type of percept could be identical to different types of neural states. Let me first discuss the issues not in terms of mind-brain identity but of correspondence. Type-type neuro-perceptual correspondences may indeed not be required logically. But consider, again, the percept of a lightness step. I would expect that whenever such percepts occur, regardless of location, orientation, size, and so on, they are always accompanied by corresponding step-shaped neural distributions at the bridge locus. This is clearly a more parsimonious proposal than the idea that lightness steps might correspond to neural distributions of several different shapes: a multiple-realizability theory would have to explain how different neural distributions come to correspond to the same percept, I am not aware of any convincing examples of neural multiple-realizability of perceptual features such as lightness, color, texture, motion, and so on. Quite the contrary; researchers tend, quite naturally, by default, to look for equal neural bases of equal types of percepts.

Finally, concerning the relation of isomorphism and mind-brain identity, the import of the statement of mine that Pessoa et al. objected to was not that isomorphism implies identity, but that identity would explain isomorphism. Suppose, for a moment, that type-type correspondences and isomorphism get wide-spread empirical support. The obvious next question would concern why such a close neuro-perceptual correspondence should hold. The lure of identity theory is that it would transform this mystery into a triviality: if two entities turn out to be identical, there is no wonder that they should be isomorphic!

The Cartesian Broadway

Christopher W. Tyler

Smith-Kettlewell Eye Research Institute, San Francisco, CA 94115.
cwt@skivs.ski.org   www.ski.org/cwt.html

Abstract: Although Pessoa, Teller & Noë make excellent points concerning the need for a mechanism of filling-in, they throw out the baby of neural specificity with the bathwater of isomorphism and the homuncular observer. The core act of perception is sensory processing by a stationary observer and does not require overt behavioral interaction with the environment. The complexity of intracortical interconnectivity does not preclude local specificity in the representation of higher-order stimulus properties.

Sparse sampling obviates the need for filling-in. Dennett (1992) is right, of course. There is no need for filling-in of stimulus regions where no seeing occurs. The classic example is in cases of sparse sampling of the image, not even mentioned in the target article. The sky looks a uniform blue, even though the "blue" photoreceptors are spaced far enough apart so that the gaps could easily be seen by the intervening bipolar cells fed by the "red" and "green" photoreceptors. But logically, we do not need to fill in the gaps if we do not survey the gaps with an "attentional" probe.

Where Dennett seems wrong, and Pessoa et al. miss the point, is in most of the cases known as "filling-in," because there does seem to be perceptual and attentional sampling within the filled-in field between the contours that generate it. The data convincingly support the existence of some kind of perceptual representation in these regions, but Pessoa et al. reject this implication as a naive isomorphism, implying that it needs a Cartesian homunculus to view the filled-in representation.

The poverty of linking propositions. Pessoa et al.'s account of the Teller & Pugh (1983) paper implies that psychophysiological parallelism is a sufficient condition for probable identity. Just because a physiological response decays at the same rate as a percept, one would not rush to assume that one explains the other. A contemporary example is provided by functional magnetic resonance imaging (fMRI) of the brain, which records properties of blood flow. The parallelism between fMRI response and perceptual properties may be extensive, but no one believes that the percept is carried by the blood. It seems that a larger set of linking propositions is required; for example, one must have a system capable of information transfer and storage before it is plausible that it could support conscious percepts.

Thus, it is easy to agree with Pessoa et al. that psychophysiological parallelism is not sufficient as a linking proposition. A major goal of their analysis is to evaluate the notion of a "bridge locus," which is essentially the "stage" in the conception of the "Cartesian theater" in which percepts are brought to consciousness. Pessoa et al. focus on the rich feedback interconnectivity throughout the cortex as a reason to doubt the existence of such a locus. As the role of such feedback may well be limited to gain control (or "auto-ranging") functions, however, it may not interfere with an essentially feedforward notion of the information processing.

What is more damning, it seems to me, is the stunning multiplicity of such representations (over 30 visual representation areas alone). The problem of identifying the homunculus in the audience of the Cartesian theater is that there are so many candidates. In one sense, each visual representation area is the audience for the Cartesian theater of the previous one. The brain is a veritable Broadway of competing shows, with each audience performing on the stage of the next theater! Is there a final arbiter or a hetarchy of interacting critics? Pessoa et al. are not explicit, proposing an explanation at the "personal" level of interaction with the environment, which makes the question of locus meaningless.

Local specificity of neural representation. To propose, however, that this level of complexity precludes any local specificity of perceptual processing, seems premature. Critics of fMRI studies often cite the complexity of the neural processing as an *a priori* reason that fMRI data will not reveal anything useful about brain processing. The data often appear to support such criticism, since they are typically obtained from observers performing some discrimination task on complex stimuli, for which a substantial number of discrete brain areas are shown to be activated.

But suppose you ask observers simply to look at well-defined stimuli and maintain a stable perceptual state? FMRI studies of responses to such stimuli, when compared with those for control stimuli differing on only one attribute such as coherent versus incoherent motion, often show activation of only a single cortical representation area. Where did the whole brain interaction and multiple representation areas go in such cases? Critics argue that they are still present, but too weak to be seen at the available signal-to-noise ratio. In the limit it is impossible to refute such an

interpretation, but to a first approximation (say, at a 10:1 signal-to-noise ratio), such whole brain interactions are insignificant in relation to the primary signal.

What does it mean for the Cartesian theater that a controlled stimulus difference may activate just one local brain area; a different one for each stimulus property? It suggests that the information for that stimulus property is stored locally, clearly eliminating the notion of a distributed representation. The mapping studies of large numbers of retinotopic visual maps already put a profound crimp in the distributed representation. Maps are themselves an explicit isomorphism. Although several of the authors cited by Pessoa et al. agree with Dennett that isomorphism is not a necessary property of perceptual encoding, both they and Pessoa et al. need to deal with the fact that isomorphism is a physiological reality at many levels of analysis: anatomical, neurophysiological, neuropathological, and the neuroimaging level.

Isomorphism may not be necessary, but it seems to be useful. The brain has evolved to have multiple isomorphisms in many sensory and motor systems. Its main utility is probably in minimizing the length of wiring while maximizing connectivity. Just as it is usually much easier to show someone how to construct a bookcase than to explain verbally, the intermediate step of an isomorphic mapping may provide a basis for manipulations of hypothetical cases that are not readily accessible to a symbolic representation. The utility of such mapping seems to have escaped critics of the Shepardian concept of mental rotation (Shepard 1981). This suggests the need for a new logical category: not necessary but efficient!

The fallacy of the personal. Briefly, Pessoa et al. do not articulate the process of perception as "personal"; they essentially assert that perception *is* the organism interacting with the environment. This viewpoint does not encompass the perception that dominates my awareness in everyday life, since it invokes a purely behavioral definition. Moreover, it would have to imply that perception is attributable to the behavior of simple organisms such as amoebae (or even automatic pool cleaners), since these show complex interactions with the environment that are superficially indistinguishable from those of humans. So the simple attribution of perception to personal behavior is insufficient to define its properties in a meaningful way.

Analytical isomorphism and Marilyn Monroe

Robert Van Gulick

Department of Philosophy and Cognitive Science, Syracuse University, Syracuse, NY 13244-1170. **rnvangul@mailbox.syr.edu**

Abstract: Pessoa, Thompson & Noë present compelling evidence in support of their central claims about the diversity of filling-in, but they embed those claims within a larger framework that rejects analytical isomorphism and uses the personal/subpersonal distinction to challenge the explanatory importance of filling-in. The latter views seem more problematic.

Pessoa et al. define the doctrine of "analytical isomorphism" as holding that "cognitive neuroscientific explanation requires the postulation of a "final stage" in the brain – a bridge locus – in which there is an isomorphism between neural activity and how things seem to the subject" (sect. 8.1, para. 2). They take issue both with the postulation of a bridge locus and the requirement of isomorphism, but their reasons for rejecting the doctrine may depend on reading it in an unnecessarily strong way. Based on Teller and Pugh (1983), they interpret the bridge locus not only as the immediate neural substrate of perception, but as a "particular set of neurons having a particular pattern of activity that is necessary and sufficient for a particular perceptual state" (sect. 8.1, para. 5). Read in this way, it does seem hard to justify the a priori requirement of a bridge locus. Why must the neural substrate reside locally in a particular set of neurons? As Pessoa et al. note, more

global patterns of activity and temporally coded assemblies provide empirically plausible alternatives. Moreover, given neural plasticity and the empirical possibility – if not likelihood – of multiple realizability, it seems unlikely that any given neural pattern is necessary in the sense of being the only possible neural substrate for a given percept.

Nonetheless, it seems that one could reinterpret the bridge locus in a way that avoids both of these problems while still providing a notion adequate for the formulation of analytical isomorphism. Assuming the truth of materialism, which Pessoa et al. do not seem to doubt, every token percept must have a neural substrate – whether local or global – that is sufficient to produce or constitute it. Given multiple realizability, the nature of that substrate may vary from one token to another, but on each occasion that the percept is achieved, there must be some neural substrate or other that is nomically sufficient for the percept. Call that state in each case the bridge locus.

Analytical isomorphism then becomes the demand that on each occasion there must be an isomorphism between percept and its bridge locus, that is, between the features of each token percept and its specific neural substrate. The issue thus shifts to the question of what counts as an isomorphism in this context. If there is to be an explanatory link between percept and substrate, there must be some way – at least in principle – to map the features of the percept onto underlying features sufficient to produce or constitute them. The absence of such a mapping would leave the link brute and inexplicable if not downright magical. To count as an isomorphism the features mapped from percept to substrate need not be the same; spatial relations in the percept need not be mapped onto spatial features in the substrate. As Pessoa et al. note, the isomorphism may be functional rather then topographic (sect. 8.1). Given this broad and more abstract notion of isomorphism, how could there fail to be an adequate mapping? The authors might reply that to avoid triviality the doctrine must place some constraints on what counts as an isomorphism. Even abstractly interpreted, the notion of isomorphism implies a relational sameness of structure. Thus, at the least, it must be possible to decompose both the precept and its substrate into elements (whether global or local) and then to establish some sort of correspondence between the relations among those elements in the two domains. Were the authors to deny that there is any meaningful decomposition of the substrate into parts or elements – no matter how broadly construed – they might argue in turn that the notion of isomorphism cannot be meaningfully applied. But I am reluctant to attribute any such view to them since they do not suggest it nor can I say it is plausible. Thus the issue remains confused. It is hard to say whether analytical isomorphism should be rejected without knowing more about what counts as an isomorphic mapping from percept to substrate. I hope the authors will clarify the matter in their reply.

Pessoa et al. draw attention to the distinction between perceptual content at the personal and the subpersonal level, a distinction which they claim is regularly ignored both by those who invoke filling in and by those who criticize it. The alleged common error is to suppose that perceiving a continuous fully detailed world involves being aware of a continuous fully detailed representation. The authors claim that, on the contrary, personal-level perceptual content is a matter of what the person (or whole organism) can accomplish in relation to the environment. "At the personal level, there are no pictures, images, or other representations in visual perception; there is simply experience of the world" (sect. 9.2). Thus, they conclude that the absence of a continuous detailed subpersonal representation is no obstacle to perceiving such a world at the personal level. Nor should the fact that our experience seems detailed and continuous in such a case imply that there is anything illusory about how things seem.

Thus, in the case of the Marilyn Monroe imaged wallpaper, Pessoa et al. attribute to Dennett (1992) the view that "the person's experience of the Marilyns is mistaken or illusory" because it seems fully detailed despite the absence of a underlying fully detailed representation of the wall. They accuse Dennett of relying on the uniformity assumption, according to which "there is no difference in kind between perceptual content at the personal level and neural content at the subpersonal level." It is difficult to believe that Dennett, who has had so much to say about the personal/subpersonal distinction, makes such an assumption. There are obviously important differences between content at different levels. Nor need one suppose otherwise to support Dennett's reasoning about the Marilyns example. What is required is not an equation of personal and subpersonal content, but merely an appeal to a plausible principle of dependence that the authors themselves endorse when they write, "we accept the general thesis that facts about brain-level content determine what the person sees or experiences." They rightly deny that dependence entails uniformity, but that leaves unspecified the sort of subpersonal representations needed to support a given form of experience.

Pessoa et al. are right that our ordinary experience is not an experience of representations, images or pictures, but rather of the world. But it does not follow that our experience of the world does not *involve* image-like representations, image-like or otherwise. One can plausibly suppose that our personal level experience of the world depends upon or is even constituted by the activation of such representations, without falling into the error of supposing that they themselves are the objects of our experience.

If so, we can reasonably ask, "what sort of subpersonal representations are needed to produce a personal-level experience of a fully detailed wall of Marilyns?" To answer the question, we need to distinguish a further third level, namely, the phenomenal level. Dennett's concern is phenomenal or conscious experience, and although all such states may be personal or whole organism perceptual states, the converse is not true. A person or organism can be perceptually informed about its environment, in the active sense preferred by Pessoa et al. without having any corresponding phenomenal or conscious experience. Thus the question to ask in Dennett's example is, "What subpersonal representations are needed to support an experiential awareness of a fully detailed wall of Marilyns?" The phenomenal realists Dennett is criticizing claim that phenomenal mental states use a distinctive medium of representation, one that is often thought of as a plenum, a determinately differentiated continuous manifold. The intent of the Marilyns case is to call that view into question; there can be no such phenomenal representation because there is no underlying subpersonal representation with the details to support it. Dennett does not mean to deny that the person has a personal level perceptual state whose content is that the wall is fully Marilyn covered. In that respect he seems much more in agreement with Pessoa et al. than the latter suppose. Moreover, what he regards as mistaken or illusory is not our experience of the wall, but the naive beliefs that we may form about the nature of our experience; it is these beliefs that he regards as the source of mistaken phenomenal realism.

The challenge to the philosophical realists – of whom I am one – is to find a way to formulate their view that can accommodate the incompleteness and gappiness of experience while still preserving the view that phenomenal states involve a distinctive mode of presentation.

Fully embodying the personal level

Francisco J. Varela[a] and Pierre Vermersch[b]

[a]*LENA CNRS UPR 640, Hopital de la Salpetrière, 75013 Paris, France;*
[b]*LPC-CNRS, Ecole Normale Superieure, 92120 Monterouge, France.*
fv@ccr.jussieu.fr   www.ccr.jussieu.fr/varela/welcome.html
pvermers@es-conseil.fr   www.es-conseil.fr/GREX/

Abstract: The target article concludes that it is essential to introduce the personal level in cognitive science. We propose to take this conclusion one step further. The personal level should consist of first-person accounts of

specific, contextualized experiences, not abstract or imagined cases. Exploring first-person accounts at their own level of detail calls for the refinements of method that can link up with neural accounts.

Pessoa et al. have built a convincing case for introducing the personal level into visual science, well beyond the neural basis of filling-in (sect. 9.2). The importance of the personal level in the target article is one of its most innovative aspects. But *how* exactly is visual experience at the personal level to be studied? We would like to take their argument one step further by underscoring here the need for a methodologically disciplined approach to first-person descriptions (Varela & Shear 1999). Only then can the personal level in visual science and cognitive science have its full impact.

Let us begin with the distinction between personal and subpersonal. This is surely a useful one, but as stated in the target article, it does not make it fully clear that the personal is situated in a contextual activity manifested primarily at the level of the subject's *practical* life. Consider the metaphor of the knife and its handle. The knife's blade is akin to the subpersonal level; it constitutes most of its surface and is where the actual cutting takes place. But the handle is what the user has access to, albeit with no direct consequences in the world. But try and use a knife without a handle. This has an important consequence: The proper level of the personal is where the subject (the cognitive user) is accustomed to handling the world in a prereflexive way, in what can be called the "natural" attitude. The consequence is that to move forward in research it is necessary to turn the subject's attention to its *own* usage, and thus to introduce a level of examination that is proper for a first-person phenomenological description.

By phenomenology we mean here a methodical quest to examine not the final result (has the cut been performed?), not merely the intentional content of mental state (what is being cut? the *noema* in technical terms), but the *mental act itself* (how do you cut when you do the cutting, or *noesis*, Varela 1996). Such an analysis will bring to the fore a number of implicit phenomenal data that remained in the periphery of attention, in the sphere of the prereflexive (Vermersch 1998). A good example is the recent study of insight in professional settings, showing how the prereflexive made explicit is the key to exploring this practical human activity (Peugeot 1999).

The significant difference between a general level of the personal/subpersonal distinction and the specificity of first-person accounts can be clearly illustrated in Pessoa et al.'s critical discussion of Dennett's Marilyn wallpaper example, which he intends as a case of active perception in which someone walks into a room and notices the Marilyns. Dennett's is a perfect example of armchair psychology: an imaginary situation that has the appearance of being concrete and specific, but in fact it isn't. Thus the proper phenomenal content gets to play no role. This criticism is actually implicit in Pessoa et al.'s claim that the *task* of vision has to be understood at the personal level; instead Dennett's analysis is motivated by considerations from the subpersonal level. Dennett means his example to be a case of active perception in which someone walks into a room and notices the Marilyns, without being concerned with a detailed scrutiny of them.

But here is where the abstractness of the exercise begins to make a difference. How is the observer instructed? What is his motivation? And what is the precise task to be accomplished? Saying: "Enter the room and tell me what you see," is not the same as saying "Go in and verify that all the designs on the wallpaper are identical under examination." Being in a dentist's waiting room and looking at the wallpaper because there is nothing else to do is still an entirely different context. Surely a request for a detailed comparison will lead to results quite different from those of passive contemplation with wandering thoughts. These, in turn, mean that all these details, the very core of Pessoa et al.'s own demand for embodiment, are not circumstantial, but essential for visual science. Thus, although we concur with the authors' critical assessment of Dennett's thought-experiment, our main complaint is

that the specific phenomenal disembodiment is not brought forth sufficiently.

We have to avoid quickly slipping into a Dennett-style heterophenomenology, as if the nonhuman/human difference did not exist. This can only be done by introducing not only the "personal as a general" condition, but also the importance of *systematic first-person descriptions* that act in *mutual constraint* with third-person descriptions (Petitot et al. 1999; Varela 1996). Pessoa et al. refer to an article (Thompson et al. 1999) in which they draw from the tradition of phenomenological psychology deriving from Husserl and Merleau-Ponty, but this connection is not made explicit in the target article. Although we can understand that they cannot grapple with everything in a single article, we do wish to underscore in this discussion the potential danger of assuming that the personal does not entail a radical change in what methodological requirements are needed to make it come alive.

To conclude: The personal level, if it is to be an active engagement in the world, needs to be explored with research strategies in which: (1) The purposes and tasks must correspond to *actual* experiences (not armchair case studies) so that the subject's context is fully taken into account and documented (the place and attitude, the materials used, the injunction proposed, etc.); (2) The description of the lived experience does not boil down to a mere final discrimination (which would be mere phenomenalism, not phenomenology), the totality of its texture (for example, the various stages of development in time) must be taken into account, which typically requires an explicit *methodological* refinement. (3) The phenomenal content of such first-person descriptions must be open to validation in mutual *circulation* with third-person studies, beyond analytic isomorphism.

Pathological completion: The blind leading the mind?

Robin Walker[a] and Jason B. Mattingley[b]

[a]Department of Psychology, Royal Holloway, University of London, Egham, Surrey, TW20 0EX, United Kingdom; [b]Department of Psychology, Monash University, Clayton, Victoria 3168, Australia. **robin.walker@rhbnc.ac.uk psyserver.pc.rhbnc.ac.uk/vision/ jason.mattingley@sci.monash.edu.au**

Abstract: The taxonomy proposed by Pessoa et al. should be extended to include "pathological" completion phenomena in patients with unilateral brain damage. Patients with visual field defects (hemianopias) may "complete" whole figures, while patients with parietal lobe damage may "complete" partial figures. We argue that the former may be consistent with the brain "filling-in" information, and the latter may be consistent with the brain ignoring the absence of information.

Pessoa et al. have provided an informative and timely review of visual completion. Missing from their suggested taxonomy, however, are the phenomena of "pathological" completion in patients with brain damage. Pathological completion was first described by Poppelreuter (1917/1990) in his pioneering studies of brain damaged patients with visual field defects. In his original study, he presented a simple shape (such as a circle) tachistoscopically so that, with central fixation, a portion of it fell into patients' impaired visual field. Under these conditions some patients paradoxically reported having "seen" a whole shape. Since these early observations, several other workers have reported similar results in patients with unilateral occipital and parietal lesions (e.g., Torjussen 1978; Warrington 1962). Although pathological completion has been considered by some to be an example of the brain "filling-in" missing information, there are a number of alternative interpretations which mirror those applied to normal completion phenomena. We have argued that there may actually be separate forms of pathological completion (Walker & Mattingley 1997;

Walker & Young 1996). One is consistent with the brain "filling-in" information, via mechanisms that are as yet poorly understood; the other appears consistent with the brain ignoring the absence of information from within a hemianopic field. Thus pathological completion encompasses both aspects of the debate surrounding normal visual completion.

Recently, we reviewed all of the studies of pathological visual completion and found that many reports of visual "filling-in" in patients with visual field defects (hemianopias) could be attributed to residual vision in the impaired hemifield, compensatory eye movements, or undetected sparing of visual function within the macular region (Walker & Mattingley 1997). When these issues are taken into account there are surprisingly few convincing demonstrations of pathological completion in patients with visual field defects. One notable exception is a study by Torjussen (1978) who tested three hemianopic individuals using flash-induced afterimages to stabilise the image on the retina (see also Weiskrantz 1990). Under these conditions, all three hemianopic patients reported seeing whole figures as complete, even though part of the stimulus fell within their objectively blind field. Moreover, half figures presented entirely within the field defect were never reported, thus reducing the likelihood that residual vision was the explanation of this form of completion. None of the patients reported seeing half-figures as complete when these were presented entirely in their intact visual field, suggesting that patients' veridical perception of whole shapes depended on part of the stimulus falling across both the sighted and "blind" visual fields. Furthermore, stimuli in the blind field were reported only when a mirror-symmetric location was stimulated in the sighted hemifield. Torjussen discounted "imaginative filling-in" (or confabulation) as an explanation for his findings. Instead, he attributed the effect to interhemispheric interactions whereby the presence of a stimulus in the sighted hemifield can facilitate functioning in the hemianopic field (cf. Pöppel & Richards 1974).

In contrast to Torjussen's findings, there are numerous descriptions of patients who report "seeing" a complete figure when an incomplete one is presented in their intact visual field. Although this seems to be an example of the brain filling-in the incomplete portion of the figure, the nature of the patient's brain damage adds a further confounding factor. We have argued that previous observations of completion of partial figures may be attributed to "*unilateral neglect*," a disorder in which affected individuals are unaware of visual stimuli falling toward the side of space opposite their damaged (typically right) hemisphere. In support of this suggestion, we (Walker & Young 1996) presented evidence from a nonhemianopic patient who had damage to the right parietal lobe and showed neglect for the left sides of individual objects, regardless of where they were located in the visual field. Our patient often reported seeing whole shapes when figures were presented with missing left sides, even when these were presented entirely within his ipsilesional visual field. In this instance apparent completion may be a further consequence of unilateral neglect whereby the brain ignores the absence of the missing left side. Furthermore, our review of the literature revealed that the overwhelming majority of cases where pathological completion of partial figures has been reported can be related to the presence of parietal neglect (Walker & Mattingley 1997).

We propose that the veridical perception of whole figures, which has been observed in a few hemianopic individuals (without neglect), and pathological completion of partial figures in patients with unilateral neglect, reflect different underlying mechanisms. Veridical perception of whole figures in hemianopic patients may be an example of the brain filling-in missing information and is open to a number of different interpretations as described by Pessoa et al. for completion in normal subjects. It may be related to completion at the blind spot, which has been attributed by some to a form of amodal completion in which the blindspot is treated as an occluder (Durgin et al. 1995). It may also reflect "top-down" activation of stored object representations, which reduce thresholds for detecting a stimulus falling in the blind field. Alternatively, it may involve interhemispheric interactions that modulate visual

sensitivity within the blind field. Pathological completion in patients with parietal neglect, however, may simply reflect their lack of awareness for the contralesional side of visual stimuli. In these cases, pathological completion may be regarded as a *consequence* of neglect, rather than a reflection of some form of active filling-in. Thus, consistent with Dennett's view on how the brain treats the blindspot, we suggest that parietal neglect patients may fail to become aware of missing contralesional information, leading to apparent completion.

Filling-in does require a mechanism, and some persistent doubts

Paul Whittle

Department of Experimental Psychology, University of Cambridge, Cambridge CB2 3EB, England. **pw109@cam.ac.uk**

Abstract: (1) In the everyday situation of seeing a uniform patch on a variegated background, filling-in does seem to require a mechanism. (2) Our attempts to make matches to such a patch make us appreciate anew the elusiveness of the phenomenon.

The target article is a useful and worthy successor to Walls's (1954) review. Filling-in is not a minor phenomenon. It is of the *essence* of vision. Vision is the only sense that presents the world to us with *no gaps.* I strongly agree with Pessoa et al.'s thesis that a notion like "active vision" is required to avoid a Cartesian impasse, but I restrict myself here to more empirical matters.

I want to discuss one basic kind of filling-in, for two reasons. First, this kind is so common, and so inescapably requires some sort of computation, that it seems to me to render academic the question of whether filling-in requires neural machinery. Second, we have been studying it, and close attention to the phenomenon is instructive.

The situation I have in mind is seeing a patch of uniform colour against a variegated background, so that the contrast varies around the edge of the patch. We usually see such a patch as uniform. But since there is good evidence from psychophysics and physiology that the colour of a uniform region depends on the contrast at its edge, in this case the various edge contrasts must be combined to yield the single colour we see. We cannot conceive how this could be done without some kind of machinery. It may be more like a balance than a computer, but machinery there must be. Even though Sherrington, as recently as 1906, found the idea of psychic combination without physiological mediation intelligible in the case of binocular fusion, this would seem a bizarre idea, to put it mildly, in the present case. (Though might it not be good for us to stretch our imaginations in a Sherringtonian direction?)

Note that there is no assumption of a "Cartesian theatre" here. The argument is that we see one colour as a result of many inputs – the various edge contrasts – so the inputs must be combined in some way. It is like the question of how a unitary colour is derived from several cone or opponent channels. My point is simply that filling-in requires a mechanism no less and no more than that does. The postulation of such a mechanism is not otiose; there is something for it to do. This does not touch the question of whether it is like painting in or finding out.

I said that we usually see such a patch as uniform. However, this uniformity can break down, and in some cases it is not achieved at all. Both are instructive, but I discuss only the first here. Thomas Hayward and I made a computer graphic of a 10° circle divided into four quadrants: white, light grey, dark grey, and black. We superimposed a 2° mid-grey "centre disk" on the centre where the quadrants met. Anyone coming into the room saw it as a uniform mid-grey disk, with no problem. We wanted to manipulate the four edge contrasts to find out what combination rule determined the exact grey. First, we tried to set up a comparison disk in a uniform surround, which subjects could adjust to match the centre disk.

The striking result was that although we tried many ways of presenting the comparison disk, we could find none that gave satisfactory matches. If it was to one side, then matches were based on the contrasts of the centre disk with the quadrants between it and the comparison disk, and were very different depending on which quadrants those were, despite the fact that the matches were all made to the same grey. If it was in the same position, but alternating slowly in time, with a blank interval between displays, then we could never make the two disks look remotely the same, even though one was mid-grey and the other could be set anywhere between black and white. It was a chalk and cheese difference, not just a zone of uncertainty.

One subject's response was that he needed another dimension: the black-white continuum was not enough. He could not abstract lightness from other characteristics of the centre disk.

My experience was that when the comparison disk was, for example, set brighter than its surround, my attention was drawn to those parts of the centre disk that were darker than their adjacent background quadrants, and they looked markedly darker than the comparison. The uniform filling-in did not survive close scrutiny. We have found the same in other quite different experiments in this lab, with patches on nonuniform backgrounds. The uniformity is a real phenomenon, but requires a global rather than an analytic perceptual attitude. "Finding-out" – attentive looking – destroys the effect! This perhaps complements demonstrations by Davis (1996) and others that filling-in of various types occurs in pre-attentive vision.

These observations suggest to us that the phenomenology of filling-in is more subtle than is often assumed. We see something striking, describe it in readily available object terms, and these are then fixed too firmly by the social process of generating a field of work, whereas we should remain open all the time to the ambiguity and elusiveness of the original phenomenon.

A task-oriented taxonomy of visual completion

Carol Yin

Computation and Neural Systems, Division of Biology, MC 139-74, California Institute of Technology, Pasadena, CA 91125. **carol@percipi.caltech.edu**

Abstract: Differences and similarities between modal and amodal completions can only be understood by considering the goals of visual completion: unity, shape, and perceptual quality. Pessoa et al. cannot reject representational accounts of vision because of flaws with isomorphic representations of perceptual quality: representations and processes for perceptual quality (modal completion) and most likely dissociable from those for unity and shape (nonmodal completions).

Pessoa, Thompson & Noë have made a much needed contribution to research on visual completion with their preliminary taxonomy of completion phenomena; I wholeheartedly agree with their call for a more task-oriented approach to vision. They have made a convincing case that there is no need to fill in an isomorphic representation to see an image. However, they then extend the flaws of isomorphic representations to all visual representations in general, and incorrectly conclude that a representational account of vision serves little purpose.

Pessoa et al.'s chain of reasoning rests implicitly on the idea that the purpose of visual completion is to provide the perceptual quality of a visible surface (i.e., brightness, color, etc.), hence the discussion about Cartesian materialism. However, visual completion has other purposes as well: determining unity and shape. These other purposes are achieved without regard to perceptual quality and do not require completion of a pictorial representation. Understanding them, however, does require a higher-order representational account of vision.

The phenomenal difference between modal and amodal completions has misled some researchers to posit completely different processes for each. Pessoa et al. have weighted this difference too lightly, and assumed that their conclusions for modal completions apply to amodal completions. The confusion can be cleared up with a better understanding of the goals of visual completion. Some goals may require different mechanisms for modal and amodal completion, but most do not. In the spirit of Pessoa et al.'s call for a more task-oriented approach to vision, I would like to propose that the taxonomy of visual completion be revised and expanded to better reflect the different goals of visual completion.

Pessoa et al.'s preliminary taxonomy assumes that there are only two tasks for visual completion: recovering shape and producing perceptual quality. The issue of figural *unity* has been neglected. The optical information that projects to the retinas is frequently fragmented or incomplete because objects may partly occlude one another. The goal of unification is to determine which fragments belong together. Rensink and Enns (1998) have shown that partly occluded fragments can be "linked" or unified preattentively, indicating the importance of unity for later visual processing. Whereas the early Gestalt psychologists were interested in figural unity (Koffka 1935), it has received only scattered attention in recent studies of adult perception (cf. Boselie & Wouterlood 1992; Gillam & McGrath 1979; Kellman & Shipley 1991; Palmer & Rock 1994; Peterson 1994; Trick & Enns 1997; Williams & Hanson 1996; Yin et al. 1997). Note that in many shape completion studies, the unity of the occluded figure is a given (Gerbino & Salmaso 1987; Sekuler & Palmer 1992; Takeichi et al. 1995).

Edge and surface-feature processes are intrinsically complementary to each other, although they may not be strictly serial. Grossberg and Mingolla (1985) first made the distinction between edge and surface-feature processes, but this distinction relegated surface features to providing only the perceptual quality of completed surface regions: Boundary processes provided shape and unity, then the filling-in of features from those boundaries provided the perceptual quality of the visible surfaces.

The revised taxonomy should reflect the contributions of surface-feature and edge processes to each of the three completion goals of unity, shape, and perceptual quality. However, the contributions of surface features to unity and shape have in the past been neglected. As Figure 1 illustrates, edge processes can provide both shape and unity, but surface features can only provide unity information and thus have an indirect role in shape. Yin et al. (1997) found evidence that surface completion processes can *amodally* unify image fragments that arise from partial-occlusion, based on the similarity of surface color and texture. Furthermore, these surface completion processes may interact with edge processes in determining shape.

The processes that are concerned with amodal versus modal perceptual quality need to be dissociated from processes that recover shape and unity. Shape processes are neither modal nor amodal[1] (Grossberg 1994; Kellman & Shipley 1991; Kellman et al. 1998; Ringach & Shapley 1996). Unity processes are also likely to be insensitive to modality for several reasons, including parsimony and the likelihood that unity is processed prior to (or at least in parallel with) shape. Perceptual quality processes, however, must be modal by definition. These processes probably stabilize only after unity and shape have been determined, since perceptual qualities such as brightness can be influenced by higher-order factors such as grouping, depth, and three-dimensional structure (Adelson 1993; Gilchrist 1977; Kanizsa 1979; Knill & Kersten 1991; Nakayama et al. 1990).

The controversy about filling in isomorphic representations is moot for the goals of unity and shape, because they can both be accomplished nonmodally, without creating or recreating a visible pictorial representation. Pessoa et al.'s arguments are valid only for the task of determining perceptual quality, and their conclusions cannot be extended to other tasks. It would be better to avoid the terms "modal" and "amodal" altogether in future discussions of visual completion and instead refer specifically to the goal of recovering shape, unity, or perceptual quality.

Surface feature completion

Figure 1 (Yin). Different contributions of edge and surface-feature processes to shape and unity: (a) Surface feature and boundary information both contribute to shape and unity; (b) Boundary information is sufficient to overcome differences in surface features, and provide shape and unity; (c) In the absence of boundary completion, some impression of unity may be provided by surface feature similarity; (d) When boundaries do not complete and surface features are dissimilar, the figure appears to be three disconnected surface fragments. Adapted from Yin et al. (1997).

We need to have internal representations that differ from what is immediately available in the optical input because we must constantly act upon things that we cannot see directly: objects may be partly occluded or may move beyond our field of view. We may not necessarily store every visual representation we derive, but storage is a separate issue. There is abundant psychophysical as well as neuropsychological evidence (see Farah 1990) that derived visual representations do underlie our behavior, even though those representations are not necessarily available to phenomenal experience. We can heed Pessoa et al.'s call to avoid positing final stages of processing by reconceptualizing visual representations as being temporally dynamic or even multiplexed. However, the future of visual completion research lies not in abandoning representational accounts of vision but in understanding how representations reflect different states of processing for different visual tasks.

NOTE

1. Anderson and Julesz (1995) point out that the horizontal arms of the cross (see Fig. 10 of the target article) may appear as bent flaps hinged to the vertical bar. This happens more frequently when the horizontal bars have near disparity. This does not entail that there are two kinds of boundary interpolation as they suggest. Rather, it shows that the L-junctions at the intersection of the cross have a bistable interpretation: a planar horizontal bar is seen when L-junctions are interpreted as T-junctions signaling occlusion, thus initiating edge interpolation (see Kellman & Shipley 1991). Bent flaps are seen when L-junctions are interpreted as L-junctions signalling corners in the boundaries of two objects that abut the vertical bar.

Authors' Response

Filling-in is for finding out

Luiz Pessoa[a], Evan Thompson[b], and Alva Noë[c]
[a]Department of Computer and Systems Engineering, Center of Technology, Federal University of Rio de Janeiro, Ilha do Fundao, Rio de Janeiro, RJ 21945-970, Brazil; [b]Department of Philosophy and Centre for Vision Research, York University, North York, Ontario, Canada M3J 1P3; [c]Department of Philosophy, University of California, Santa Cruz, Santa Cruz, CA 95064. **pessoa@cos.ufrj.br   www.cos.ufrj.br/~pessoa**
evant@yorku.ca www.yorku.ca/research/vision/evant.htm
anoe@cats.ucsc.edu www2.ucsc.edu/people/anoe/

Abstract: The following points are discussed in response to the commentaries: (1) A taxonomy of perceptual completion phenomena should rely on both phenomenological and mechanistic criteria. (2) Certain forms of perceptual completion are caused by topographically organized neural processes – neural filling-in. (3) The bridge locus, understood as the final site of perceptual experience in the brain, should be replaced by the principle that each token percept has a neural substrate that is nomically sufficient for it, all else being equal. (4) Analytic isomorphism – the view that there must be a pictorial or spatial neural-perceptual isomorphism at the bridge locus – should be rejected. Although more abstract kinds of isomorphism are central to the neural-perceptual mapping, the perceptual cannot be exhaustively explained in terms of the neural, and therefore the explanation of perception cannot be reduced to uncovering neural-perceptual isomorphisms. (5) The task of vision is to guide action in the world, not to construct a detailed world-model in the head. (6) Neural filling-in facilitates the integration of information and thereby helps the animal find out about its environment. (7) Perceptual content needs to be understood at the level of the person or animal interacting in the world.

The main aim of our target article was to examine issues surrounding the topic of visual filling-in, and to provoke discussion from visual scientists and philosophers. Our discussion focused on the conceptual and methodological issues of neural-perceptual isomorphism, the bridge locus, representation and action, and the personal/subpersonal distinction. These issues, though theoretical, influence everyday empirical research. For example, **Walker & Mattingley**'s studies of pathological completion (which, unfortunately, we did not know about when we wrote our article) have at their core one of the main issues we discuss: Is perceptual completion "ignoring an absence" (as seems to be the case for neglect patients) or "representing a presence" (as seems to be the case for some hemianopic patients)? Moreover, the issues are not confined to visual science, but pertain to all sensory modalities, as **Appelbaum** clearly demonstrates in her discussion of analytic isomorphism in speech perception. We are pleased that a majority of the commentators address these issues, and we wish to thank them for their mainly constructive contributions.

As we stated in the target article, the term "filling-in" is used in so many different ways that confusion can easily arise unless researchers are careful about what they mean. Some of the commentaries perpetuate such confusion. **Kranda** states that we favor "the traditional notion of filling-in," despite our clear distinction between perceptual completion and neural filling-in, our rejection of the traditional isomorphic framework, and our review of the cases in which there is evidence for neural filling-in and the cases in which there is none. At the other extreme, **Tyler** inter-

prets us as rejecting neural filling-in, although we took two sections (6 and 7) to discuss the evidence for it. Kranda and Tyler should read **O'Regan**'s commentary, for he recognizes that although we think there is evidence for neural filling-in – evidence with which he takes issue (see sect. 2) – we also think that neural filling-in needs to be severed from the traditional notion of the filling-in of an inner pictorial representation that is isomorphic to our perceptual experience at the bridge locus.

Our response falls into six sections (see Table R1). Section R1 is devoted to taxonomic issues. Section R2 restates the case for neural filling-in. Section R3 discusses the bridge locus and neural-perceptual isomorphism. Section R4 clarifies our criticism of the "representationist" conception of vision. Section R5 shows how to reconcile what some commentators (**Durgin, Kaufman, O'Regan**) perceived to be two incompatible aspects of our position – our argument for neural filling-in, and our argument that vision is an active process whose goal is not representation (in the sense of constructing a world-model in the head). Finally, Section R6 discusses perceptual content and the personal/subpersonal distinction.

R1. Taxonomic issues

R1.1. Phenomenological *and* mechanistic criteria. Perceptual completion phenomena are diverse and encompass different visual modalities. We argued in the target article that both phenomenological and mechanistic considerations should guide the elaboration of a taxonomy, and we proposed that the modal/amodal and boundary/feature axes be used to classify completion phenomena.

Consider the modal/amodal distinction, which **Dresp** and **Yin** argue is not useful, because it is phenomenological rather than psychophysical (Dresp) or goal-oriented (Yin). Modal completion produces visible attributes, such as color and brightness, whereas amodal completion does not. **Yin** writes that the phenomenological differences between modal and amodal completion have misled researchers into supposing different mechanisms for each; she argues that we have underestimated these differences and assume that "[our] conclusions for modal completions apply to amodal completions." Her interpretation of our position is mistaken, however, for we clearly state in the target article that further experimental research is needed to determine the

similarities and differences between these two types of completion. As an example, consider the work of **Davis & Driver,** which supports our view of the perceptual importance of the modal/amodal distinction. As they discuss, small, local stimulus manipulations, such as a change in disparity sign, can change completion from amodal to modal (see their Fig. 1, and Nakayama & Shimojo 1990a). In the case of amodal completion, visual attributes at occluded (and completed) regions do not reflect properties of the occluded object, but of the occluder. In modal completion, the occluding (and completed) region inherits attributes from adjacent areas. Why should this be the case? Davis & Driver suggest that visual attributes help the visual system draw attention to modally completed regions. Hence they propose a difference in the functional effects of the two types of completion that may be associated with the visual system's attending to potentially relevant information in one case and ignoring irrelevant information in the other.

Contrary to **Dresp**'s claim that our taxonomy is purely phenomenological, we invoked the modal/amodal distinction not primarily because of the perceptual differences between these two forms of completion, but because we wished to pose precisely the question of whether the two forms of completion engage similar mechanisms. The issue is not whether a purely mechanistic taxonomy is superior to a phenomenological one (as she suggests), or vice-versa, for both kinds are needed for a comprehensive understanding of visual perception in general and perceptual completion in particular. The issue is rather how the types of taxonomy are related to each other. For example, do modal and amodal completion share boundary interpolation mechanisms? Some evidence favors a common source (Kellman & Shipley 1991), whereas other studies indicate separate processes (Anderson 1995). Even if the two kinds of completion do engage similar mechanisms, the distinction between them at the perceptual level would not be invalidated, for as we argued in the target article, to eliminate the phenomenological dimension (as Dresp advocates) is to lose sight of vision as, in the first instance, a personal-level phenomenon (see sect. 6).

R1.2. Area, surface, and contour completion versus boundary/featural completion. A new taxonomy that includes area, surface, and contour completion is proposed by **Dresp.** How do these types of completion relate to our boundary/feature axis? Contour completion is equivalent to

boundary completion. We view surface completion as the complex interplay between boundary and featural completion processes; in other words, surfaces emerge as basic units of visual perception as the result of the interactions between these two complementary completion systems (Grossberg & Mingolla 1985). This view is based on experimental studies that indicate the existence of two separate completion systems (see sect. 3.3.3, and De Weerd et al. 1998). We thus agree with **Yin** that both boundary and featural processes contribute to surface determination (shape and unity).

What about area completion? Given that **Dresp** advocates a purely mechanistic taxonomy (in which there are only as many forms of completion as there are distinct completion mechanisms in the brain), should she distinguish between area and surface completion? The finding that "areas" show decreasing detection thresholds from the border (Dresp & Bonnett 1993), whereas "surfaces" do not (Dresp 1992), suggests that the associated experiments are tapping into distinct levels of the visual system. It does not suggest distinct mechanisms, however. We submit that Dresp's area completion and surface completion rely on common neural mechanisms, namely, the spreading of contrast signals that occur at luminance borders and are regulated by contour signals. For example, one of us has described a brightness model that can account for both what Dresp calls area completion (such as Mach bands) and the quality (or feature) part of surface completion (Pessoa 1996a; 1996b; Pessoa et al. 1995). Thus Dresp herself seems to rely on a distinction between phenomenological and mechanistic criteria of classification.

R1.3. Unification: Is it grouping or completion? As advocates of a task-oriented approach to vision, we welcome **Yin**'s suggestion that the taxonomy be expanded to reflect the different goals of perceptual completion. She proposes three main goals of perceptual completion: the determination of unity, shape, and perceptual quality. Shape determination is closely related to boundary completion, whereas perceptual quality determination is closely related to featural completion. Unification – the determination of which scenic fragments belong together – was largely ignored in the target article.

When considering unification, it is crucial to assess the differences between grouping and completion. Grouping is a central process of early vision; it "places a non-visual link between those fragments thought to belong to the same object in the world" (Rensink & Enns 1998). For example, in a field of randomly moving dots, a group of dots sharing the same motion may be experienced as unified. The key here is that although the dots compose a perceptual unit, no new visual elements are posited. The question that arises, then, is whether unification processes are just grouping processes, or whether some kind of completion is involved.

Rensink and Enns's recent study sheds light on this question. They used a visual search task where target and distractors included, for example, long and short bar segments, together with cube-like objects, in different geometric arrangements (suggesting, for example, occlusion of the cube by bar segments). They showed that rapid completion is highly sensitive to the existence of gaps, with linking occurring only across a completely occluded space. Such sensitivity is not characteristic of rapid grouping in

which gaps may be ignored. They find that "occlusion edges are removed and that occluders are needed to trigger linking, suggest[ing] that rapid completion may involve occlusion-specific mechanisms not found in any general grouping process."

In conclusion, we embrace **Yin**'s proposal to include unity as one of the goals of completion, and we agree that the taxonomy "should reflect the contributions of surface-feature and edge processes to each of the three completion goals of unity, shape, and perceptual quality." Care must be taken, however, to distinguish between unification processes that engage general grouping mechanisms and those that engage true mechanisms of completion.

R1.4. Pathological completion. One omission from our taxonomy is pathological completion as discussed by **Walker & Mattingley.** Such phenomena occur in patients with unilateral brain damage who report seeing whole shapes even though large portions of the stimulus fall in the impaired visual field. In their recent comprehensive review, Walker & Mattingley (1997) describe patients with visual field defects (hemianopias) that complete whole figures (with portions in the impaired field), and patients with parietal lobe damage that complete partial figures (in the intact visual field). Their analysis shows that whereas the former case is consistent with the brain filling-in, the latter is consistent with the brain ignoring the absence of information. In fact, the latter case may be attributed to unilateral neglect, a syndrome in which individuals are unaware of visual stimuli that fall within the space opposite their damaged hemisphere. Walker and Mattingley suggest that the apparent completion in neglect patients may have to do with their failing to become aware of the missing information. In this case, subjects would be ignoring the absence of information.

R2. The case for neural filling-in

R2.1. What do we mean by "neural filling-in"? In our target article (sects. 5.1, 6, and 7) we distinguished between two senses of "neural filling-in." The first is the weak notion of the brain's merely providing a signal (representing a presence as opposed to ignoring an absence), whereas the second is the strong notion of the brain's providing a signal that is spatially or topographically organized. The right conclusion to draw from our review of the empirical literature is that there is evidence for the weak kind of neural filling-in in some cases (e.g., Murakami's [1995] motion aftereffect investigation), and for the strong kind of neural filling-in in other cases (e.g., illusory contours in the studies by von der Heydt and colleagues, and brightness perception in the studies by Paradiso and colleagues discussed in the target article). Hence we disagree with **O'Regan** when he concludes that there is no evidence (at all!) for the strong kind of neural filling-in. Although there is no *proof*, the *evidence* is strong. On the face of it, the evidence is not merely for dynamic long-range interactions in the sense of interactions involving nonclassical receptive fields (see **Spillmann & Werner,** and Spillmann & Werner 1996), but for dynamic long-range interactions that involve the spreading of spatially organized signals. It is always possible, of course, to propose alternative interpretations of psychophysical results and cellular responses, but we think our interpretation is faithful to the data we cite.

On the other hand, we certainly share **O'Regan**'s concern not to "prejudge the issue of an internal, pictorial representation" (and to that extent we agree that the term "neural filling-in" can be misleading and that a better term might be sought). Indeed, one of the main tasks of our article was to criticize analytic isomorphism, which assumes that there *must* be such an inner pictorial representation. Whether there are such representations is a matter to be decided empirically. We therefore disagree with **Murakami** that "there is no promising way to test whether there is a pictorial neural representation of filling-in that is directly or indirectly equated with our perception." Although we certainly do not think that any inner representation (pictorial or otherwise) is to be *equated* with our perception (see sect. 6), whether there are (subpersonal) pictorial neural representations causally responsible for personal-level perceptual experience must be decided experimentally. (In the target article we mentioned a number of promising recent experiments that bear on this issue, such as De Weerd et al. [1995] on texture filling-in and Rossi et al. [1996] on the cortical substrates of brightness perception.)

R2.2. What neural mechanisms mediate perceptual completion? It is proposed by **Spillmann & Werner** that three types of neural processes may be responsible for perceptual completion phenomena: (1) converging feedforward projections (from lower to higher levels) leading to the generation of new receptive field properties at higher levels; (2) horizontal intrinsic connections allowing interactions within a single visual area; and (3) feedback projections to cells at lower levels that respond selectively to specific properties of a stimulus (see also Spillmann & Werner 1996). Figure R1 shows these three types of mechanisms. For example, responses to luminance borders may be propagated to inner regions (that do not receive direct support) by horizontal connections and thus form the neural substrate for the Craik-O'Brien-Cornsweet effect.

Although mechanisms of these three kinds may be able to account for perceptual completion, at the present moment not a single completion phenomenon can be traced directly to a particular type of neural mechanism. Consider illusory contours. Although it is well established that the responses of certain cells in primate V2 correlate well with illusory contour stimuli, how these cell properties come about is unknown. Indeed, computational models belonging to each of the three classes above can be found in the literature, such as the feedforward model of Peterhans et al. (1986), the horizontal-connections-based model of Neumann et al. (1997), and the feedback model of Grossberg & Mingolla (1985).

Bachmann adds a fourth kind of possible completion mechanism – the intriguing notion that modulatory inputs from nonspecific arousal systems, such as nonspecific thalamic nuclei, could play a role in neural filling-in. These nuclei would provide slower, spatially spreading modulation that would allow for completion. This suggestion is motivated by Bachmann's "requirement of maximum veridicality." Nevertheless, neural filling-in that is mediated by (among other things) lateral connections (which Bachmann rejects) may play a crucial role in removing ambiguity from local stimuli and hence in sustaining "adequate and consistent responding to the full concreteness of the environmental stimuli."

In sum, our current knowledge of the neural substrates of perceptual completion is sketchy at best. A concerted effort to integrate physiological recordings, psychophysical experiments, and computational modeling is needed to mitigate our ignorance. In this connection, the recent interest shown by physiologists in filling-in is encouraging (see **De Weerd, Komatsu,** and **Murakami**).

R2.3. Cortical neurons and surface perception. It is well known that uniform areas typically do not elicit vigorous responses from cortical neurons; spatiotemporal contrast is needed. Hence one of the great puzzles of cortical physiology is how neural responses encode extended surfaces. One possibility, which is suggested by the spreading type of neural filling-in, is a topographic encoding based on the activation of neurons in early cortical areas, but this type of substrate has proven elusive (Komatsu, personal communication; von der Heydt et al. 1996). For this reason, **Komatsu** proposes that surfaces are encoded by a population of size-sensitive cells across several cortical areas, with the larger receptive fields of higher areas helping to encode larger surfaces. (It is worth noting that in this scheme the percept is correlated with activity that is distributed across several visual areas, not with a single one as the traditional notion of the bridge locus would imply. We take up the bridge locus issues in sect. R3.1) **Arrington** extends this notion and proposes that "edge junctions" in general provide local surface constraints that are bound together and provide the substrate for surface encoding. These proposals notwithstanding, the question of how the visual system encodes surfaces


Figure R1.   Potential neural substrates of filling-in. (A) feedforward convergence. (B) horizontal connections. (C) feedback connections.

is far from settled. In fact, there is evidence that, in some cases, cells in the topographically organized striate cortex do respond when the entire receptive field is uniformly stimulated (Komatsu et al. 1996; Rossi et al. 1996).

R2.4. The blind spot. At the present time, the question of how the brain treats the receptorless region of the retina – whether it ignores it or actively completes it – cannot be answered. Contrary to **Kranda,** however, the issue is not a pseudo-problem and must be settled experimentally. Furthermore, in-principle arguments about the apparent lack of utility of neural completion **(Pishva)** miss the point that the underlying circuit is there for other important reasons (see below). As **Gregory** aptly remarks, once a complicated completion machine is invented, why not use it?

Most studies in primates have shown that the representation of the blind spot in V1 is monocular and driven by the ipsilateral eye – that is, by the eye whose stimulation is not at the optic disc (Horton 1984; Le Vay et al. 1985). To a first approximation, then, the associated singularity in cortical representation would provide an automatic mechanism for completion – the blind spot is simply sewn up, as it were. A bar across the blind spot would have a continuous cortical representation (although shortened in length relative to other cortical regions of comparable eccentricity), as would texture and color. Such an idea is favored by many researchers, including **Kranda** and **Neumann,** but the situation is not likely to be so simple, for the study by Fiorani et al. (1992) shows that cells can be driven *binocularly* at the cortical region corresponding to the blind spot. Does this result tell us that the brain does in fact actively fill in the blind spot? Such a conclusion would be premature for the reasons discussed in section 4.5 of the target article: The monkeys were anesthetized (see also sect. R2.6 below) and only bar stimuli were used; furthermore, the neural responses need to be compared more directly to psychophysical performance on an appropriate task.

These points show that **Durgin** is mistaken when he implies that we interpret the Fiorani et al. (1992) study as showing that "completion neurons" are "painting lines into conscious experience." Once again, the neurons reported by Fiorani et al. are important insofar as previous studies mapping the cortical representation of the blind spot reported it as being purely monocular (e.g., Horton 1984). Although the Fiorani et al. study does not allow us to conclude that a neural substrate for blind spot completion has been identified, it does raise the possibility that the neural substrate involves more than ignoring an absence or "implicit filling-in" as proposed by **Neumann** (abstract).

The same possibility is raised by **Murakami** and his study of the motion aftereffect after adaptation to filled-in motion at the blind spot (Murakami 1995), which we discussed in section 7.1 of the target article. **Durgin** charges that our argument based on this study – an argument that Murakami himself gives in his commentary – is "misleading," and suggests that his recent studies of texture aftereffects (Durgin & Cole 1997) show that we need to be more careful in the interpretation of completion data. We are aware of the dangers of inflated explanations based on particular studies. Nevertheless – and this was our main point in the target article – Murakami's study shows that real motion and filled-in motion share a common early pathway, and it provides a measurable effect of one kind of perceptual filling-in of the blind spot; therefore, his study force-

fully challenges the "neglect" or "ignoring an absence" hypothesis about blind spot completion.

To date the most telling findings about the cortical representation of the blind spot come from psychophysical investigations. One important study by Tripathy et al. (1995) also shows why **Kranda**'s putative "paradox" about filling-in mechanisms does not withstand scrutiny. They asked subjects to compare the lengths of test bars presented at the blind spot to the lengths of reference bars presented at the corresponding location in the other eye. They found either small distortions or no distortions at all in perceived length at the blind spot. More importantly, performance at the blind spot was comparable to that in normal periphery at a comparable retinal location (eccentricity). How do these results help us understand the cortical representation of the blind spot? If opposite sides of the blind spot are contiguously represented in the cortex, then size distortions should occur in the task, because in this case the effective length of the bar is its total physical length minus the diameter of the blind spot in question (the part that is not physically stimulated). What occurred, however, were only very minor size distortions, which suggests that the blind spot is completed with the background (by whatever mechanism), rather than being sewn up. Recent psychophysical results such as these (see also Tripathy & Levi 1994) suggest that blind spot completion is not a simple matter of a singularity in cortical representation, though they do not specify how completion is accomplished.

R2.5. Perceptual completion without cortical reorganization. In section 4.3 of the target article, we discussed studies of cortical reorganization after retinal lesion and the idea that the associated receptive field expansion may help explain the perceptual filling-in of color and texture, as well as the appearance of illusory contours (see Gilbert 1992). **Murakami** mentions a recent study of his that questions this idea (Murakami et al. 1997). He and his colleagues showed (1) that there is behavioral evidence for perceptual filling-in at a scotoma produced by a *monocular* retinal lesion; and (2) that no cortical reorganization occurs in monkey area V1. Instead, the normal visual system seems to be able to produce perceptual filling-in when some part of the retina is lesioned. (Of course, there remains the possibility of perceptual filling-in resulting from cortical reorganization in the case of *binocular* lesions as suggested by Gilbert [1992]).

R2.6. Awareness and cortical cell responses. In section 4.5 of the target article, we urged caution when one compares perceptual results with the results of single-cell recordings in anesthetized animals. **De Weerd** provides additional material relevant to this point. He describes unpublished experiments with Ricardo Gattass on texture filling-in with anesthetized Cebus monkeys in which they did not find "climbing activity" in V2 and V3 under conditions that caused both perceptual filling-in in human observers and climbing activity in neurons of awake Rhesus monkeys. Although this finding might be specific to Cebus monkeys, it is more likely that the state of the animal was the determining factor in obtaining a negative result. As De Weerd aptly states: "The absence of climbing activity during unconsciousness enhances its possible role as a neural substrate for perceptual filling-in." Should these results withstand careful scrutiny, they reinforce our point in the

target article about how to interpret results in anesthetized animals. In this context it should be noted that the results of Fiorani et al. (1992) in anesthetized Cebus monkeys (discussed in sect. 4.2 of the target article) have recently been replicated in anesthetized skunks (Leticia & Volchan, unpublished observations), a species with a simpler brain than that of primates.

R2.7. The time-course of filling-in and normal surface perception. One key task in linking the perceptual and neural domains is to relate the time-course of perceptual completion to that of underlying neural mechanisms, a matter we discussed in section 4.5 of the target article. **Dennett** notes that effects that take longer than around 200 msecs of fixation to develop, such as the texture filling-in effects studied by Ramachandran and colleagues (Ramachandran & Gregory 1991; Ramachandran et al. 1993), cannot themselves be part of normal vision. They can tell us about normal vision, however, as **De Weerd** helpfully elaborates in his commentary by addressing the issue of why different types of featural completion (e.g., texture and brightness) follow different time courses. He and his colleagues have shown in an impressive recent study (De Weerd et al. 1998) that the slow time scale of texture filling-in (several seconds) does not reflect a slow filling-in process, but rather the long time it takes for figure-ground segregation to fail, after which a fast completion process ensues. In other words, after the contours between the textured and untextured regions adapt (see Fig. 13 of the target article), they can no longer contain the fast spreading process that results. In everyday perception contours antagonize the perceptual completion of texture information, but weak boundaries can be overcome and produce fast perceptual completion. As De Weerd et al. (1998) state: "[t]he seconds-long delay before perceptual filling-in illustrates the close intertwining of segregation and neural filling-in processes, which is at the basis of normal surface perception." Finally, these results further corroborate the dichotomy between boundary and feature systems advanced by Grossberg and Mingolla (1985), and hence provide additional support to the boundary/feature axis in our perceptual completion taxonomy.

R2.8. The level of filling-in. Although the exact mechanisms of neural filling-in are unknown, what we do know suggests that they occur early in the process of vision. For example, the perceptual completion of boundaries in illusory contours occurs as early as V2 (see sect. 4.1 and 6.1 of the target article). In fact, all the effects suggestive of neural filling-in that we reviewed are indicative of mechanisms in V1, V2, and V3.

Gregory asks whether "illusory surfaces" (as perceived in a Kanizsa figure) also occur early. He observes that "completion seems to depend on general rules rather than on specific object knowledge" and suggests that such rules can be considered to be cognitive (see also **O'Regan**). We think, however, that the kind of rules exhibited in surface completion are probably not readily susceptible to the subject's beliefs and desires, and so are probably not cognitive in this sense. One of the hallmarks of cognition is its flexibility. Yet the stereotyped character of completion – for example, its dependence on the length, number, and contrast of inducers (see Lesher 1995) – marks it as typical of "early vision."

What do psychophysical results reveal about this issue? **Dresp** suggests that the processes that determine illusory surface brightness occur at higher levels in the visual system than the ones responsible for "local" brightness effects, such as seen in Mach bands (see also Dresp 1992). This suggestion implies that they occur at least beyond the striate cortex (see Pessoa 1996a), which does not rule out that much. Davis and Driver (1994) showed that detection of Kanizsa figures in a visual search paradigm is parallel. Hence, the perception of illusory figures does not require focal attention, and occurs at early stages of the visual system. Unfortunately, Davis and Driver's results are not conclusive with regard to surface completion, for they are consistent with the filling-in of boundaries between the inducers (as we discuss in the target article, boundary and surface completion may involve different processes). Their results show that boundary completion is preattentive, but not that surface completion is preattentive. What needs to be devised, therefore, are direct tests for preattentive surface completion. One would need stimuli that produce illusory surfaces but no contours. Although one would have to be careful about controls, one could use stimuli such as those in Figure 8A of the target article. Such experiments would enable us to answer more directly **Gregory**'s question about the level of illusory surfaces.

A recent study by Mattingley et al. (1997) also lends support to the idea that early factors are involved in surface completion. They showed that parietal extinction arises only after the formation of visual surfaces, such as Kanizsa figures (see especially their Experiment 4). Although we cannot review their study in detail here, their results provide neurophysiological evidence that surface filling-in occurs early.

R2.9. Neural filling-in, prediction, and statistical inference. It is suggested by **Barlow** that prediction should be seen as pattern completion in the temporal domain, as filling-in in time (see also **Sauvé**). Both spatial filling-in (as discussed in the target article) and temporal filling-in (as discussed by Barlow) can be seen as types of statistical inference. Viewing perceptual completion as statistical inference and prediction could be of help to task-level descriptions of vision. Barlow goes on to state, however, that "statistical inference, jumping to a conclusion, subjective filling-in, and spatially propagating neural activity are surely different levels of description of the same process in the brain and there need be no incompatibility between them." Were we wrong, then, in our criticism of Dennett (sects. 6 and 7 of the target article), when we assumed that spatially propagating activity is more than "jumping to a conclusion"? We do not think so. In some cases the processes might be one and the same, though described at different levels, but in other cases they could indeed be different. For example, in texture completion as studied by Watanabe and Cavanagh (1991; 1993), our guess is that the visual is just "jumping to a conclusion" in Dennett's sense, that is, simply signaling that "more of the same" is present. No topographic representation or neural spreading through lateral interactions (or interactions of some other form) would be involved in this case, whereas they do seem to be involved in illusory contours and brightness perception. Precisely because there are so many different levels of description for completion phenomena, it is easy to slide unwittingly from one to the other. Our view is that careful

experimental work should determine our descriptions of the phenomena.

R3. The bridge locus and neural-perceptual isomorphism

In the target article (sect. 8.1) we criticized the doctrine of analytic isomorphism, according to which there must be a final stage in the brain – a bridge locus – where an isomorphism obtains between neural activity and the subject's perceptual experience. We took issue with both the postulation of a bridge locus and the *a priori* preference for isomorphism in the neural-perceptual mapping. Although we focused on vision, we meant our points to apply to the study of perception in general. We therefore wish to thank **Appelbaum,** who shows how analytic isomorphism comes close to being a dogma in the study of speech perception, one which, she argues, "will have to be rejected before progress in explaining speech perception will be made."

Todorović and **Van Gulick,** on the other hand, think that our reasons for rejecting analytic isomorphism depend, as Van Gulick puts it, "upon reading it in an unnecessarily strong way." They defend a more refined, nonlocalizationist conception of the bridge locus – as does **Hahn** – and then suggest that isomorphism is the most plausible way to understand the neural-perceptual mapping. In this section we will consider their arguments and the related points made by **Barlow, Lehar, Myin, Pishva, Ross,** and **Tyler.**

R3.1. The bridge locus. We endorse **Barlow**'s claim that it is useful for neuroscientists to seek "the location in a neural system where one can most readily show that particular neural states provide enough information to account for particular psychological performance." Let us call such neural states the informational substrate of the performance or percept. Several commentators – **Barlow, Hahn, Todorović,** and **Van Gulick** – seem to think that we were attacking this conception of the informational substrate, but that was not our intention at all. This conception is *not* what Teller and Pugh (1983) meant by the term "bridge locus," contrary to Barlow's claim, nor is it what we criticized in the target article. As we read Teller and Pugh, the bridge locus is the set of neurons whose activation is necessary and sufficient for a percept.[2] Barlow's informational substrate is not sufficient to produce a percept, although it is necessary (all other things being equal) for the production of a percept. Indeed, as Teller has advised us (personal communication), Barlow's informational substrate is what she and Pugh called the "critical locus":

> Most visual scientists also would accept [in addition to the bridge locus as the immediate substrate of perceptual experience] the notion that in a particular experiment the flow of information (e.g., about the spectral content of a stimulus) through the visual system can be importantly blocked, limited or otherwise reorganized at a particular stage. . . If one stage imposes such a critical transformation, this stage will be called the *critical locus.* . . for that experiment (Teller & Pugh 1983, p. 581).

Far from wishing to challenge the idea of the critical locus, we feel that it is a useful one for visual science. We agree with **Barlow** who writes: "States of activity in underlying neural representations *must* provide sufficient information for the perceptual discriminations that are based on

them." He continues by noting that the discovery of such neural representations "enables one to find out when brain events can and when they cannot explain perceptual phenomena."

We doubt, however, that the stronger conception of the bridge locus is useful for shaping research in visual science. This was the central claim made in our discussion of the concept in the target article. We reject the notion of the bridge locus for the following reasons. First, there is no logical requirement that we postulate it. Second, the empirical evidence to date does not lend strong support for the view that there must be a *final* neural stage in processing: (1) Brain regions are not independent stages or modules; they interact reciprocally because of dense forward and backward projections, as well as reciprocal cross-connections. (2) A central idea in recent neuroscience is that brain activity is highly distributed (spatially and temporally). (3) Neural plasticity in recovery from lesions suggests that the connection between a specific psychological performance and a particular locus is variable and context-dependent. In general, therefore, we think that there is little empirical support for the idea that vision is a process whereby retinal information streams to a final place "where consciousness happens."

In the target article, we assumed that the requirement of spatial and temporal locality was built into the bridge locus concept, and we therefore argued that the concept should be abandoned. **Hahn, Todorović,** and **Van Gulick** comment that the concept can be revised so that it no longer incorporates this assumption. (Davida Teller – one of the original formulators of the concept – also recently made this point [Teller 1998].) Van Gulick's formulation is the clearest and most precise:

> Assuming the truth of materialism . . . every token percept must have a neural substrate – whether local or global – that is sufficient to produce or constitute it. Given multiple realizability, the nature of that substrate may vary from one token to another, but on each occasion that the percept is achieved, there must be some neural substrate or other that is nomically sufficient for the percept. Call that state in each case the bridge locus.

The basic idea, then, is that the core of the bridge locus concept is not the localizationist assumption, but rather the assumption that each percept has a neural substrate that is *nomically sufficient* for it.

Subject to two essential qualifications, we are prepared to accept this proposal. First, the neural substrate is sufficient to *produce* the percept, but not to *constitute* it. Perception is an adaptive, cognitive capacity of the animal or person, and therefore it is constituted not by the neural substrate *per se,* but by the interactions of the nervous system, the body, and the environment (see Chiel & Beer 1997). Second, the neural substrate is sufficient to produce the percept *only given the assumption that all other things are equal.* The "all other things being equal" clause is crucial, for the neural substrate is nomically sufficient for the percept only if the following three conditions exist: (1) the neural substrate is properly embedded in the nervous system; (2) the nervous system is properly embodied in the animal or person; and (3) the animal or person is properly situated in the environment. The absence of any of these three conditions would provide a reason for holding that the percept would not be present, even if its typical neural substrate were. Our commitment to (3) in particular puts us –

to use current philosophical jargon – in the "externalist" camp: Perceptual content is inherently world-involving, and therefore cannot be reductively identified with neural states inside the individual.

Ross, on the other hand, suggests that instead of thinking of a perceptual bridge locus, we should instead consider "behavioral bridge loci." These are neural activity patterns that must be reached to produce motor actions. The main issue is how the visual system updates its activity in the face of current stimuli so as to enable visually guided behavior. This idea is certainly congenial to our enactive, task-level approach to vision. In the case of perceptual completion, the issue would be how the topographically organized neural processes found in some cases help to support and enable behavior. We take up this topic in section R5.

In sum, if we use the term "bridge locus" to mean the neural substrate that is nomically sufficient, all other things being equal, to produce a given token percept, then we have no objection to the concept. Its status, however, strikes us as similar to that of the filling-in concept: When properly clarified and revised it seems a shadow of its former self. Indeed, just as one might argue that the very term "neural filling-in" can be misleading (see our earlier response to **O'Regan** in sect. R2.1), so one might argue that the very term "bridge locus" can be misleading, because it does not make explicit the above qualifications (besides having the lingering localizationist and Cartesian materialist connotation of a final stage in the brain).[3]

R3.2. Neural-perceptual isomorphism. In the target article, we criticized what we call "analytic isomorphism," which is the idea that there *must* be a neural-perceptual isomorphism at the bridge locus. **Van Gulick** suggests that our rejection of this doctrine depends on taking the requirement of isomorphism in an unnecessarily strong way. He argues that "[i]f there is to be an explanatory link between percept and substrate, there must be some way – at least in principle – to map the features of the percept onto underlying features sufficient to produce . . . them." all that is needed for isomorphism, he suggests, is that such a mapping exist (see also **Arrington**). He points out that "[t]o count as an isomorphism the features mapped from percept to substrate need not be the same; spatial relations in the percept need not be mapped onto spatial features in the substrate." Thus his idea seems to be that it is sufficient for neural-perceptual isomorphism that we can in fact *explain* elements or facts at the level of the percept in terms of corresponding processes or patterns of activation at the neural level. The absence of this kind of "broad and more abstract" isomorphism, he argues, would leave the neural-perceptual link "brute and inexplicable if not downright magical."

In the target article, it was not our intention to attack this highly abstract and fairly weak notion of isomorphism. Our aim, rather, was to criticize the doctrine that there must be a pictorial (spatial) isomorphism at the bridge locus. The idea is that the neural substrate will be isomorphic to the corresponding percept in the sense that the spatially organized properties of the percept will correspond to the spatially organized properties of the neural substrate. Such an equality of spatial structure would provide a kind of "transparent relation" between the neural and the perceptual domains, to use **Myin**'s term. It was this notion of pictorial isomorphism that we criticized in our discussion of **Todorović**'s treatment of the Craik-O'Brien-Cornsweet effect.

We are willing to endorse **Van Gulick**'s "broad and more abstract" neural-perceptual isomorphism, but only subject to an important qualification. As stated above (sect. R3.1), and in the target article (sect. 9), we think that there are likely to be significant limitations on the extent to which the perceptual can be explained solely in terms of the neural. A central point of our target article was to call into question the view that explanations of the perceptual can be given entirely in neural terms without including references to the environment in which the animal or person is embedded, as well as nonneural facts about the animal's body (see also sects. R4 and R6 below). To employ a different philosophical jargon, we believe that perceptual states do not supervene on neural states, but on neural and environmental states, and it is therefore not possible to explain *every* feature of perception in neural terms. For example, it is an open empirical question to what extent the categorical structure of human color perception can be explained in neural terms without including linguistic and cultural factors (see Hardin & Maffi 1997 and Saunders & van Brakel 1997 for discussion of this issue).

Because, in general, it is not possible to give an exhaustive explanation of the perceptual in terms of the neural, the explanation of perception cannot be reduced to uncovering neural-perceptual isomorphisms. Thus we think that the doctrine of analytic isomorphism, even as modified by **Van Gulick,** is a misleading way of thinking about the relation between the personal and subpersonal bases of perception. Nevertheless, we think Van Gulick is right that, to the extent that it is possible to explain the perceptual in terms of the neural, such explanation will proceed by establishing a mapping between the perceptual and neural domains, and this mapping will typically be highly abstract and functional, rather than pictorial or topographic.

Myin helpfully draws attention to an important feature of our position, which is that we reject only the *necessity* of neural-perceptual pictorial isomorphism, not its *possibility.* Whether there are such neural-perceptual isomorphisms in any given case strikes us as an empirical matter for cognitive neuroscience to decide (see sect. 8.1 of the target article). Thus there is nothing about our position that requires us to oppose **Tyler**'s correct observation that the visual system does in fact make use of retinotopic maps. **Todorović** states that he is happy to transform the idea of a (topographic) neural-perceptual isomorphism (in the case of brightness perception) from an *a priori* assumption into an ordinary hypothesis. Contrary to Tyler's reading of our position, this transformation is exactly the move we recommend for visual science. Furthermore, we agree with Todorović that there is nothing objectionable in the *working hypothesis* that there are pictorial isomorphisms at the bridge locus (when the bridge locus concept is watered down in the way we just outlined in sect. R3.1).

As we discussed in the target article, Köhler (1947) argued that neural-perceptual isomorphisms hold between the dynamic-topological relations in each domain, not the topographic or geometric ones. The problems with his field-theory view of brain activity notwithstanding, we believe that such a dynamical system approach to isomorphism can provide a more promising avenue for research in cognitive neuroscience than searching for spatial isomorphisms. Indeed, this emphasis on the dynamic-topological features of the neural and perceptual domains has been revitalized by the so-called dynamical approach to cognition,

which also incorporates elements from the Gibsonian ecological perspective (see Kelso 1995; Port & Van Gelder 1995).

Pishva offers another proposal. He suggests that we abandon the *a priori* demand for an isomorphism (of any kind) between the neural and the perceptual domains, and instead take the neural-perceptual mapping to be *homomorphic,* that is, some of the structural properties of the perceptual domain could map onto structural properties of the neural domain, without there being a one-to-one correspondence between the relations among the elements in the two domains. This is an interesting suggestion, but Pishva provides no clear example of what such a mapping would look like in visual science, and so his proposal is difficult to evaluate.[4]

In attacking the analytical requirement of pictorial isomorphism are we attacking a paper tiger? We think not. For example, **Lehar** explicitly endorses it: "[I]somorphism suggests that the required algorithm [for space perception] involves the construction of a fully spatial three-dimensional internal model of that external world as observed subjectively." One could not ask for a better example of *a priori* computational theorizing, which in this particular case happens to fly in the face of what we know about the neural substrates of space perception: Space does not seem to be represented in a uniform three dimensional map separate from the processes involved in motor actions and selective attention, but rather in a multiplicity of highly specialized visuomotor maps (see Milner & Goodale 1995; Rizzolatti et al. 1994).

Although this kind of neural-perceptual isomorphism may be an easy target, even theorists who take pictorial isomorphism to be just a working hypothesis are sometimes so wedded to the idea that it becomes a dogma. In such cases, the quest for pictorial isomorphism can hamper research by preventing other possibilities from being explored. The history of Mach band investigation provides a case in point. Mach's brilliant insight allowed him to propose that some form of lateral-inhibition circuit mediated the effect that now carries his name. Lateral inhibition *per se,* however, cannot explain Mach bands, insofar as they are not elicited at sharp luminance transitions (or edges), but occur only when there are gradual luminance transitions, such as in luminance ramps (Pessoa 1996a). The lateral-inhibition account predicts that Mach bands should be stronger at edges. So why was such a glaring discrepancy not noticed for more than a hundred years? We think that the strikingly close spatial resemblance between the percept and the retinal cell responses obtained by Hartline and colleagues (Hartline 1940; 1949) led researchers simply to ignore this key fact. The spatial isomorphism account of Mach bands seemed so natural that the absence of such bands at edges was dismissed as less relevant. It was only when the retina became less favored as the crucial site of the effect and more sophisticated accounts of visual function emerged (e.g., the postulation of "bar" and "edge" detectors), that accounts less centered on the "spatial resemblance" principle appeared (see Pessoa 1996a). Once that occurred, explaining why ramps produce Mach bands and why edges do not became central to understanding the underlying circuitry.

One final point. The main empirical moral of our target article in the case of perceptual completion is that there is no *a priori* or empirical reason to insist on only one kind of isomorphic or nonisomorphic neural-perceptual mapping.

Thus we think that in the case of texture completion the evidence does not favor pictorial isomorphism (or strong neural filling-in), whereas in the case of brightness perception it does. Nevertheless, given the misleading connotations of the pictorial isomorphism concept (the tendency to take it in the analytic or Cartesian materialist way), we think it would be much more informative and precise to talk about *topographically organized* neural substrates, and to discuss the properties of such computational maps for perception and action (e.g., Knudsen et al. 1987; Schwartz 1980).

R4. Representation and the task of vision

In section 9.1 of the target article, we argued in favor of a task-based conception of vision at the level of the animal or person interacting with the world, and contrasted this "enactive" conception with the "representational conception of vision." Even though a large part of our article was devoted to considerations about the content of various subpersonal neural processes, our critical points about a certain kind of representational model of vision led a number of the commentators – **Byrne, Davis & Driver, Earle, Lomas, Myin, Singh & Hoffman, Van Gulick,** and **Yin** – to interpret us as rejecting the very idea of subpersonal representation in vision. This was not our intention. What we meant to criticize, rather, was the idea that perception is representational in the sense of requiring a world-model in the head. To avoid further misunderstanding, we would like to clarify our position here.

First, to understand how vision works one must understand what vision is, namely, a cognitive capacity of the animal, the function of which is to give the animal access to the environment and to guide action, not to produce a unified internal representation of the outside world from the retinal image. Therefore, it is a (conceptual) mistake to answer the question "what is vision?" by saying that vision is the recovery of the outside world in a model inside the head after the fashion of inverse optics, though such an assertion is of course perfectly legitimate as a hypothesis about how vision works subpersonally. We are well aware that Marr's (1982) "work was characterized by careful analyses of the tasks to be solved by particular visual processes" (**Davis & Driver**; see also **Earle**). The point of our critical remark about Marr in the target article (sect. 9.1) was that he and his followers (e.g., Poggio et al. 1985) often treat the subpersonal task of the visual system (which they define as the ill-posed inverse optics problem of deriving object-centered representations of the world from viewer-centered two-dimensional retinal images) as if it were the cognitive function of vision at the level of the animal or person. It is this tendency that can unwittingly lead to the fallacious form of homuncularism, or to the mistaken idea that the distinction between the animal or person and its subpersonal states does not matter in cognitive science (e.g., Fodor 1975, p. 52). In contrast, Gibson's (1979) attention to the animal as the proper subject of perception enables us to avoid these problems by providing a proper description of what vision has to accomplish at the level of the animal (Nakayama 1994; Noë 1995, Ch. 4; Thompson 1995, pp. 232–35).

Second, the issue is not merely one of how to understand the concept of vision, for the proper characterization of the

task of vision at the level of the animal can lead to better scientific accounts of how vision works subpersonally. For example, if one starts from the abstraction of the stationary retinal image and asks how the visual system manages to derive a model of the objective world, one will decompose the visual system into modules that are passive in the sense of not being interconnected to motor processes. On the other hand, if one starts from the sensorimotor cycle of saccadic eye movement and asks how the animal is able to fixate points in the environment, one will decompose the visual system into visuomotor modules that guide action. As it turns out, such an animal-centered, action-based decomposition reduces the need for certain kinds of representations in vision, in particular for an on-line, moment-to-moment, detailed world-model (see Ballard 1991; 1996), a result that happens to cohere with psychological experiments on scene recognition (see Intraub 1997; Simons & Levin 1997) and neuropsychology (Milner & Goodale 1995).

We trust these remarks demonstrate that we do not think there is any "opposition" (**Singh & Hoffman**) or "inconsistency" (**Earle**) between the personal-level and the subpersonal-level in accounts of vision, and that we do not "create a gulf between neural and personal accounts of perception that belies the role that neural accounts play in explanations of visual experience on the personal level" (**Lomas**). (This latter accusation ignores sects. 4, 6, and 7 of the target article, which are devoted to the role that neural accounts play in explanations of perceptual completion.) What we object to is the neglect of the personal level and the way it motivates the faulty subpersonal theory of vision as the passive reconstruction of the world in a model inside the head. The reason this theory is defective is that it focuses on the kind of representation that is not useful or even necessary for the animal. For example, constructing a detailed representation in world-centered coordinates is unnecessary and hence a poor strategy if actions can rely on representations in a more convenient gaze-direction coordinate system (Ballard 1991; 1996).

The connection between this point and the topic of perceptual completion is that if one thinks that vision does require a detailed world-model, then one will take neural filling-in to be obligatory for vision (see **Lehar**). On the other hand, to reject the representation-as-world-model conception forces us to reconsider the role played by neural filling-in. Although there is evidence for topographically organized processes of neural completion, there is no reason to think the function of these processes is to fill-in the detail of a single, unified representation or world-model in the head. What, then, is the function of these processes? A full answer to this question would take another article, but in the next section we would like to sketch an answer based on the idea that filling-in is for finding out, and that finding out depends on the integration of local signals.

R5. Filling-in is for finding out

R5.1. When is filling-in needed? Active filling-in mechanisms are postulated by **Grossberg,** first, to overcome the inadequacies of the initial retinal coding of the image (e.g., occlusion by blood vessels), and second, for the reconstruction of surface features. Neural responses to uniform regions are largely suppressed at the retinal level, and thus to reconstruct visible surfaces one needs to rely on the neural responses generated at object borders as contrast estimates (see Neumann & Pessoa 1998a and Pessoa & Neumann 1998 for an explicit link between diffusive filling-in mechanisms and reconstruction through regularization theory). Yet Grossberg does not explain why such reconstruction is needed, aside from saying that "without further processing, visual representations would be colored outline cartoons" (see sect. 9.1 of the target article for a critical discussion of this idea). In fact, cell responses at borders carry enough information to determine appearance unambiguously in regions where responses were originally suppressed. For example, a brightness step can be signaled by, say, retinal On-cell responses at the light side and Off-cell responses at the dark side. Although the final filling-in computations in models such as Grossberg and Todorović's (1988) help determine how well the model's predictions match perception, they do not add information to what was already present in border-related responses (brightness is proportional to the integral of On-responses minus Off-responses). The final diffusive filling-in step could have been omitted.

Grossberg also states that without filling-in it is not possible to explain why regions with equivalent early neural signals (such as responses at the retina) have different brightnesses, or why regions with different neural signals have the same brightness. Consider Figure R2. The one-dimensional luminance and brightness profiles of a step stimulus are sketched together with a schematic representation of retinal responses (above baseline responses correspond to On-cell responses, and below baseline responses to Off-cell responses). In Figure R2(A), two regions with identical (zero) responses are associated with different brightnesses (see circles), whereas in Figure R2(B), two regions with different neural signals have the same brightness (see circles). Although neural filling-in is one mechanism to obtain the correct response, a host of alternative accounts have been proposed in the literature (see Pessoa 1996a). These accounts typically include as their central element the detection of primitive "events," such as lines and edges, and possibly some "symbolic" encoding. (The key feature of all accounts is using contrast information at luminance transitions to specify brightness levels.) What careful consideration of Figure R2 reveals, therefore, is that the processes that determine brightness perception must make use of *nonlocal* information, because the cells that fire for such stimuli are at or around the luminance discontinuities (at the center).

R5.2. Using filling-in for finding out. How can the topographically organized processes suggested by some of the filling-in data (e.g., brightness and illusory contours) facilitate the visual behavior of the animal in its environment? We need to address this question because several commentators (**Bachmann, Durgin, Kaufman, O'Regan,** and **Ross**) see a missing link in our target article between neural filling-in data and task-level considerations about vision. We believe that there are important task-level or behavioral reasons for neural filling-in, and that as a result, the evidence we provide in support of neural filling-in is compatible with a task-based, personal-level approach to vision. In a nutshell, neural filling-in helps the animal find out about its environment.

To link neural filling-in to the situated visual behavior of the animal in a satisfying way we would need, on the one

Figure R2. Luminance to brightness mapping (one-dimensional cross sections). A luminance step or edge produces a cusp-shaped distribution of retinal activity. The final percept comprises uniform brightness levels arranged as a brightness step. Depending on the spatial location, the same retinal responses may be associated with different brightness levels (A), or distinct retinal responses may be associated with the same brightness level (B).

hand, a more careful description of the processing tasks the visual system faces, and on the other hand, a deeper understanding on a case-by-case basis of the ways in which neural filling-in can meet those processing demands. In the case of brightness perception, for example, what are the reasons that the visual needs to go beyond the representation at retinal ganglion cell responses, or cortical simple and complex cell responses (localized edge responses)? To answer this question, we need to determine the usefulness of the *results* of filling-in (or integration) computations.

To understand the function of neural filling-in, consider that vision has evolved to guide action. For action, animals need reliable information about their environment. The visual system in its earlier stages, however, acquires only reliable *local* estimates (as a result of the spatially confined profiles of receptive fields in early vision). These estimates are ambiguous with respect to the larger global context. Therefore, to fulfill its action-guiding role the visual system must integrate this available information. There are a number of different ways in which integration could be accomplished in principle. We believe that neural filling-in is one way for the visual system to accomplish this task of integrating early information (which must be useful for finding out about the environment; see also **Tyler**). Although neural filling-in is not necessary, it may be an efficient integration method.

Consider the particular case of lightness perception. At the retinal level, cells encode stimulus contrast. Ratio sensitivity enables the visual system to handle adequately, within the limited firing range of visual neurons, the more than five orders of magnitude over which the incoming illumination varies. The dynamic range of such cells is thus greatly expanded. Nevertheless, sensitivity to luminance ratios is not enough to provide reliable estimates of surface properties, such as lightness and color, because local esti-

mates can be completely ambiguous. For example, consider a sequence of abutting patches of ascending luminance values (forming a luminance staircase), such that at every transition (or step) physical contrast is defined to be identical. The early neural responses to a transition at, for example, the beginning of the entire stimulus are identical to responses near the opposite extreme. At the beginning of the staircase, a dark gray patch is seen, whereas at the end a much lighter one is perceived. The only way that a reliable estimate of surface appearance can be obtained is by *combining* the information of all scenic contrast-responses of interest.

Lightness perception, therefore, requires the integration of information (see also **Whittle**).[5] How might this be accomplished in the visual system? One way is simply to rely on retinal contrast information and consider the entire array of stimulation (i.e., neither fill-in nor integrate). Let us refer to this operation as *implicit integration.* Now consider the case where a simple creature uses lightness information to guide its behavior – for example, to flee from "light-gray objects."[6] In this case, the entire set of optic fibers would somehow have to guide behavior. (Remember that local signals are ambiguous with respect to lightness values.) The alternative would be to use the initial retinal firing pattern to produce a more compact neural representation through integration of retinal contrast measures – which we could call *explicit integration.* This transformation would provide more direct accessibility to lightness signals and facilitate perceptuo-motor communication (see Fig. R3). In the process, something like a "lightness image" would be generated in the brain. There is nothing mysterious about such an image; it is there simply to support effective visually guided behavior. (We hope that these points, as well those below, convince **Yin** that we do not view the purpose of completion as providing visual surface perceptual quality.)

Figure R3. Sensory-motor mapping. (A) Early sensory responses to image contrast produce distributed spatial patterns of activation at an intermediate processing stage. For robust determination of scenic attributes (such as contours, lightness, and color) at a given location (black disk), all responses must be used by motor processes (implicit integration). (B) Early sensory processing responses are explicitly integrated and produce a localized response for that spatial location.

Although the above example suggests that integration is needed, how such integration should be accomplished is an open question. For completeness, we considered the possibility that some form of implicit integration occurs. Such a hypothesis, that the entire array of contrast responses constitutes by itself the underlying lightness signal, is superficially attractive, given that nothing extra needs to be computed. The problem, however, is that such a proposal could be iterated indefinitely without any criterion for stopping (see **Lehar**). Thus it could be claimed that the pattern of cone responses is sufficient to explain the percept, which is clearly unsatisfactory. On the other hand, the argument applies the other way around, too. Some criterion must be established to determine when the neural states provide enough information (**Barlow**). In other words, how much integration is enough? We submit that in many cases this question cannot be answered from a purely perceptual perspective; the behavioral side of the equation must also be entered.

Similar types of argument for the need for integration apply to other forms of perceptual completion, such as boundary completion in illusory contour perception. In general, contours are not specified solely by luminance discontinuities (or edges). Several other scenic attributes, such as texture, depth, and motion, contribute to the specification of *object* contours. As in the case of lightness, local signals are insufficient because they are highly ambiguous. For example, in the case of an object contour occluded at several regions by mud splashes, there is very little information about the whole contour locally at the unoccluded parts (Grossberg et al. 1997); also see **Shipley** for a related example). Nevertheless, in many such cases, the perception of a single contour is effortless. As in the case of lightness, integration allows for stable perception, in this case of object contours. Boundary-formation processes are important, as Shipley observes, because occlusion is prevalent in the world. Such processes, then, might aid in the identification of objects seen moving through a cluttered world.

In sum, we submit that the integration of information is needed for robust perception and action. Such integration removes ambiguity and creates the basis of more robust, in-

variant representations. As **Sauvé** observes, it increases the continuity of objects and features across space, improving the accessibility of the relevant information. We believe that integration is what is revealed by some filling-in experiments: Filling-in is caught in the act of integrating. Notice that we say that *integration* is needed, not filling-in.

In light of the evidence for neural filling-in in topographically organized regions, we think it would be useful for visual science to investigate further what advantages such integration schemes have over other nonspatially organized methods. One reason might be related to the fact that cortical topography minimizes the total volume of axons and dendrites (Cowey 1979; Mitchison 1995). Hence the local computations supported by topographic maps might favor the integration of information that preserves spatial structure, and in this way effective use would be made of the cortical circuitry. In such cases, filling-in would be accompanied by what seem like spreading interactions and spatially organized representations.

R6. Perceptual content and the personal/subpersonal distinction. In section 9.2 of the target article, we argued that the evaluation of perceptual content should not be dictated by subpersonal considerations about internal processing, but rather by considerations about the task of vision at the level of the person or animal interacting with the world. We contrasted this view with the "uniformity of content thesis," which maintains that perceptual content at the personal level is just neural content at the subpersonal level. Contrary to **Hahn,** our rejection of this thesis has nothing to do with an "endorsement of dualism," but rather with trying to make sense of the different kinds of content found in the natural world. We hold that perceptual content pertains to the animal as an embodied agent interacting with the environment, requires for its description an animal-centered, task-level account, and is constrained by norms of rationality. Conversely, neural content pertains to the animal's internal functional organization and requires for its description levels of explanation concerned with internal processing. We are grateful to the commentators who addressed this topic, for their remarks enable us to clarify, revise, and improve our position.

R6.1. The uniformity of content thesis: Dennett and visual science. Our interpretation of **Dennett** and our criticism of the uniformity of content thesis are contested by **Byrne** and **Van Gulick.** The interpretative questions about Dennett are important to the extent that they touch on substantive issues about perception at the personal and subpersonal levels. Rather than respond to them point by point, however, we think it will be more fruitful to reformulate our position in light of the commentaries by Byrne, Dennett, and Van Gulick.

One of our targets in section 9.2 was a certain pattern of reasoning that relies on the uniformity thesis. According to this reasoning, for a person to have a perceptual experience of environmental detail, the environmental detail must be represented in his brain. Therefore, if one can show that the detail is not represented in the person's brain, it will follow that the person does not really see all the detail. We drew attention to this idea in the target article in connection with Dennett's "Marilyns" example. The reasoning in that context goes like this: From the supposed fact that the brain does not represent each of the Marilyns on the wall, it fol-

lows that the person does not really see them all. This pattern of reasoning clearly relies on the uniformity thesis.

We have argued that the central task of vision is not to produce detailed internal representations, but to guide action (see sect. R4 above and sect. 9.1 of the target article). If one accepts this point, it then becomes an empirical question whether the visual system must produce detailed internal representations to guide action. It is also an open empirical question whether the person's perceptual awareness of the detailed wall of Marilyns requires that there be subpersonal neural representations (pictorial or otherwise) of all the detail. The line of reasoning that we criticize treats the absence of a detailed world-model as a sufficient reason for denying that the person sees the detail. To reason in this way is to assume the uniformity of content thesis and the associated representationist conception of vision.

This pattern of reasoning, far from being a "straw man" (**Hahn**), is widespread in visual science, especially in the current studies of scene perception and change blindness (which we discussed in sect. 9.2; see also Intraub 1997; Noë et al., forthcoming; Simons & Levin 1997). The moral of these studies is generally taken to be not only that the visual system lacks a precise representation of the visual world, but also that because it lacks such a representation the visual world is an illusion (see Blackmore et al. 1995; Intraub 1997; O'Regan et al. 1996). **Durgin** relies on this reasoning in his commentary when he invokes the so-called "Grand Illusion of perception – that the perceived world is complete and fully detailed in all directions, despite the unarguable evidence that the internal representation of the world that we can study in visual cortex is nonuniformly sampled (i.e., blurry and poorly localized in the periphery)."

We believe that visual perception probably does not involve detailed pictorial or symbolic neural representations of a scene (a world-model). We also think, however, that this fact has no immediate bearing on the evaluation of the content of the person's perceptual states. In particular, we do not think it implies that our visual experience of the world is an illusion (see Noë et al., forthcoming). The reason is that we disagree with the uniformity of content thesis. In our view, one cannot read perceptual content directly from neural content, because perceptual content is constituted by the coupled systems of animal and environment, not neural states.

The assumption that perceptual content can be specified solely by virtue of neural content leads to conceptual and phenomenological distortions. Both kinds of distortion seem to be at work in **Dennett**'s discussion of filling-in (despite **Van Gulick**'s point that "[i]t is difficult to believe that Dennett, who has had so much to say about the personal/subpersonal distinction, makes such an assumption").[7] In the Marilyn wallpaper example, Dennett, after asserting that the brain just "jumps to the conclusion that the rest are Marilyns," begins a new paragraph by saying: "Of course it does not seem that way to you. It seems to you as if you are actually seeing hundreds of identical Marilyns" (1991, p. 355). Dennett does not deny that, in one sense, the person does see the Marilyns on the wall, for he continues: "And in one sense you are: [T]here are, indeed, hundreds of Marilyns out there on the wall, and you're seeing them." Yet a natural reading of these three sentences suggests that he thinks there is *another sense* in which the person does not actually see the Marilyns. What would this sense be? What Dennett goes on to say is this: "What is not

the case, however, is that there are hundreds of identical Marilyns represented in your brain." We agree with this statement, taken as an empirical claim about the person's brain. Yet the remarks give the impression that there is a sense of "seeing" or "having a visual perceptual experience," such that the person does not in fact perceptually experience hundreds of identical Marilyns – even if it seems to him that he does – because there is no detailed representation of them – pictorial or otherwise – in his brain. (Or, to put the point the other way round, if all the detail were represented in his brain, then he would in fact have a perceptual experience of the detail.) In other words, contrary to **Byrne** and **Van Gulick,** the person suffers from a kind of perceptual illusion: It seems to him that he visually experiences the Marilyns, but he does not, because the subpersonal neural representation required for such an experience is not present. Although the passage does not logically entail this idea, its tone and rhetoric strongly suggest it. (The same sort of idea is strongly implied by Dennett's remark, quoted in the target article, that the continuity of consciousness at the personal level is actually an illusion, because the underlying neural contents are discontinuous. The idea goes beyond any straightforward empirical claim, for it amounts to a view about the conditions of application at the personal level for the concept of seeing or of having a visual perceptual experience: What the person can be said to experience perceptually is the content of what his brain represents. This is a subpersonally motivated sense of seeing: It implies that, even if it seems to the person that he sees the detail, and he correctly judges that the detail is present there on the wall, he nonetheless does not in fact perceptually experience the detail unless it is neurally represented (pictorially or otherwise). This line of reasoning is unacceptable because we have reason to believe both that the detail need not be neurally represented (there is no need for such a world-model in the head) and that the person sees the detail (he says he sees it and the judgment he makes on the basis of his visual experience is correct). Because the uniformity of content thesis is responsible for this faulty line of reasoning, we should reject the thesis.[8]

Van Gulick asks: "What sort of subpersonal representations are needed to produce a personal-level experience of a fully detailed wall of Marilyns?" This is a legitimate question. We believe that we do have such experiences and that they do not require having neural representations of all the detail (symbolic or pictorial world-models) in our heads. Van Gulick observes that not all personal-level perceptual states are conscious (though all conscious perceptual states are personal-level phenomena). We accept this point, too. When we suggest, however, that there is no sense in which the person does not see the Marilyns simply by virtue of not having a neural representation of all the detail, we do not mean that the person perceives the wall unconsciously; we mean his perceptual experience of the Marilyns on the wall is a conscious one. Why do we believe that we can be conscious of the detail even though we do not represent all the detail in a scene-model in our heads? The reason is that vision is not a "snapshot" phenomenon; vision is a temporally extended process of looking constituted through eye-, head-, and whole-body-movements. As **O'Regan** (1992, p. 484) puts it: "[S]eeing constitutes an active process of probing the external environment as though it were a continuously available external memory." In other words, we do not need to represent the detail of the world in our heads be-

cause we can visit it with our eyes (see also sect. R4.1 above). The person does not visually discriminate all the Marilyns at once, but as an active moving perceiver he can explore the scene. In that sense, the person does see the detail (though this will depend on the circumstances and the task at hand: see **Varela & Vermersch**). To appreciate what is involved in consciously seeing the wall, we must reflect more carefully on the actual character of looking than most philosophers and visual scientists have been inclined to do (see Thompson et al. 1998). In other words, we need a careful conceptual and phenomenological study of perceptual experience at the personal level. (For our views on what such a study does and does not involve see our response to Varela & Vermersch below.)

In the target article, we took Dennett's remarks to imply that he believes that the person's perceptual experience is mistaken or illusory in the Marilyn case. **Byrne** counters that the mistake "is *not* any kind of perceptual illusion," but rather an "intellectual error we make when reflecting on our experience," namely, the mistaken belief that because we experience the wall as covered with detailed Marilyns, the neural representation must contain detailed representations of Marilyns. **Van Gulick,** too, writes: "what [Dennett] regards as mistaken or illusory is not our experience of the wall, but the naive beliefs that we may form about the nature of our experience." Byrne makes the same point about its seeming to one that the Marilyns are present in one's experience or mind, not just on the wall: The error is the intellectual one of invoking "something like sense data as the evidence that must be presented to an 'inner Witness' in the brain." Similarly, Van Gulick takes Dennett to be criticizing the philosophical doctrine of "phenomenal realism," which holds that "phenomenal mental states use a distinctive medium of representation, one that is often thought of as a plenum, a determinately differentiated continuous manifold."

As we indicated two paragraphs earlier, however, **Dennett** does give the impression that there is a kind of perceptual illusion involved in the Marilyn case. It is possible that he would not endorse the line of reasoning that gives rise to this idea. In any case, he certainly thinks that ordinary, unreflective perceptual experience typically involves a commitment to false beliefs about its own nature in relation to the brain. Thus he believes that in the Marilyn case the person's perceptual experience typically involves the false belief that the mind or brain must contain a detailed picture of the Marilyns. **Van Gulick** is therefore closer to the truth than **Byrne,** because he recognizes that the mistake is supposed to be one of "naive belief," rather than "intellectual error." The intellectual error is analytic isomorphism, which is the naive belief spelled out in the form of a theory. In other words, for Dennett, the mistaken intellectual doctrines of analytic isomorphism and Cartesian materialism have their origins in the naive and false beliefs of ordinary perceptual experience (rather than, as we believe, in a particular sort of misguided representationalism in philosophy of mind and cognitive science).

Dennett gives voice to this very idea in his commentary. He writes: "What needs accounting for is not necessarily that there is filling-in but that there seems to be filling-in – subjects (naively) believe there is filling-in [parentheses in original]." Then later, describing the Marilyn case he says: "I say there is something illusory about the ordinary perceptual experience, not because I endorse the isomorphism

principle, but because ordinary people do tacitly endorse it, and it leads them to error. . . It is, if you like, a theorist's illusion, but it turns out that we are all theorists." We took issue with exactly these points in the penultimate paragraph of section 9.1 of the target article. We deny that subjects naively believe there is neural filling-in, that is, we deny that their ordinary perceptual experience (or even their perceptual experience of illusory figures and their own blind spot) typically involves beliefs in neural filling-in. Ordinary perceptual experience typically involves beliefs about the world, but not beliefs (whether naive or intellectual) about what goes on in the brain when we perceive. Nor does tacit endorsement of analytic isomorphism play a role in everyday perceptual experience. To suppose that perceptual experience involves such beliefs is to falsely intellectualize perception. Thus we think that Dennett gets wrong the epistemic commitments implicit in ordinary perceptual experience. To get these commitments right requires a careful conceptual and phenomenological account of perceptual experience at the personal level.[9] Our contention is that such an account would make no mention of beliefs about neural filling-in, or of doctrines about the neural-perceptual relation such as analytic isomorphism and Cartesian materialism; instead, it would describe how people experience the world given their perceptual capacities and conceptual resources.

Finally, we wish to address the relation of the uniformity thesis to analytic isomorphism. Analytic isomorphism is not entailed by the uniformity thesis. It is a much stronger version of the requirement of uniformity of content. **Byrne** points out that many of our arguments against **Dennett** are arguments against the claim that to see the wall of Marilyns there must be a picture-like, isomorphic representation in the brain; and he argues that none of these arguments against analytic isomorphism carries any weight against the uniformity thesis. We regret the lack of clarity in our presentation that gave rise to this criticism. (Several referees objected to an earlier version of sect. 9.2 because they found it too "philosophical." In our revisions we inadvertently sacrificed clarity for brevity.) What we hope now to have made clear is that (1) we give good reasons to reject the uniformity thesis; (2) we give good reasons to reject analytic isomorphism; (3) we do not assume that arguments against analytic isomorphism count against the uniformity thesis; and (4) there is some evidence that Dennett relies on the uniformity thesis. We do, however, concede to Dennett, and to Byrne, that Dennett does not rely on analytic isomorphism.

R6.2. Consciousness and the personal level. A number of comments are provided by **Revonsuo** about "subjective visual awareness," of which only a few speak directly to points we made in the target article. He believes that "visual awareness only refers to *one* specific phenomenon," and asserts that "it is difficult to deny" that this specific phenomenon "somehow resides in the brain." He then falsely states that we "seem to deny this." We made no such denial. If by the phrase "subjective visual awareness resides in the brain," Revonsuo means that neural activity is nomically sufficient (all other things being equal) for visual experience, then we agree. What we did say was that the bearer of perceptual states is not the brain or nervous system *per se,* but rather the animal or person; and we argued that we must be careful to avoid unwarranted *a priori* assumptions

(e.g., about neural-perceptual isomorphism) when investigating the neural perceptual relation. It is not clear whether Revonsuo objects to these statements. He does say that if we mean to claim that "the *phenomenal* content of visual awareness" exists only at the personal level, then we are "obviously wrong," because it "is brought about by ('subpersonal') causal processes entirely confined inside the brain, as the generation of dream phenomenology shows." This argument is question-begging and misses the point. To say that neural processes are causally responsible for consciousness does not mean that consciousness should be conceptualized as a subpersonal-level phenomenon, rather than a personal-level one. The example of visual experience in dreams is no counterexample, for although neural activity of a particular sort is causally responsible for dreaming, it is the animal or person who dreams.

Revonsuo also says that we "seem to be unable to make the fundamental distinction between the phenomenal model of the world, constructed by the brain, and the physical world itself." This comment shows that he has not grasped the most basic points of our target article. We argued that one must clearly distinguish *what is represented* from the *medium or vehicle of representation*. In other words, one must not assume that the brain models the world by constructing a neural model that has the same properties as the world it models. We also argued that there is no reason to *assume* that the brain constructs detailed models of the world. Revonsuo does make this assumption: The phenomenal consciousness of visual perception "constitutes the brain's real-time *model* of the world, which we in everyday thinking naively take as the external world itself, as if physical reality were somehow 'directly' perceived" (see also **Lehar**). It is difficult to know how to assess such a vague, sweeping claim. We have already discussed the reasons for doubting that the brain is in the business of constructing detailed models of the world (sect. R4), and what is wrong with the idea that ordinary perceptual experience is committed to naive beliefs about its own nature in relation to the brain (sect. R6.1). In the target article we took no position on whether perception is "direct," but the issue is not nearly as obvious as Revonsuo seems to think, for it turns on the precise senses given to the terms "direct" and "indirect," and on the precise specification of the items that might or might not mediate perception.

Earle and **Tyler** raise questions about the relationship between the perceptual experience of the person and the ecological, task-based aspect of the personal level. Earle says that from an ecological, task-based perspective "the conscious experience of the perceptual world . . . can be regarded as epiphenomenal," and therefore our approach "reduces the importance of the phenomenological experience of vision." He believes that "[t]he phenomenology of vision may be important in its own right . . . but it has little to do with current visual science." Tyler's assessment of our task-based approach is similar – he says that we "essentially assert that perception is the organism interacting with the environment" – and objects because our "viewpoint does not encompass the perception that dominates my awareness in everyday life, since it invokes a purely behavioral definition." Our position, however, is that consciousness cannot be conceptually severed from the embodied life of the animal or person embedded in the world. Therefore, to study the perceptually guided action of the animal *is* to study the way in which it is aware of the world (including itself), and hence its sentient or conscious life. This view hardly treats consciousness as epiphenomenal. Nor does it reduce seeing to behaving; it simply insists on not separating perceptual experience from embodied action. Such an approach does not entail – as Tyler suggests – that "simple organisms such as amoebae (or even automatic pool cleaners)" have perceptual experience, for such organisms or systems do not (*pace* Tyler) "show complex interactions that are superficially indistinguishable from those of humans." They are easily distinguishable from those of humans.

Singh & Hoffman and **Varela & Vermersch** raise general questions about methodology at the personal level. Singh & Hoffman accuse us of not clearly distinguishing between perception as experienced from the first-person perspective and perception as studied from the third-person perspective of the scientist, and they ask whether we think the same methodology would apply in both cases. Varela & Vermersch ask: "*How* exactly is visual experience at the personal level to be studied?" They also assert the importance of disciplined "first-person phenomenological description." These important issues go beyond what we can discuss extensively here, but we do wish to make a few points.

First, we find the dichotomy between first-person and third-person perspectives problematic. When one is asked to reflect on and describe one's visual perceptual experience, one is being asked, in the first instance, to describe what one sees, that is, how the world looks, not a private inner sensation or feeling. (One can proceed to describe the perceptual act itself, but such description needs to proceed on the basis of the prior account of what one sees.) Furthermore, the descriptions one gives are produced in an intersubjective, linguistic context.[10] Therefore, we need to move beyond first-person versus third-person formulations. Second, we do not find **Varela & Vermersch**'s analogy between a knife and its handle useful for conceptualizing the relation between the subpersonal and personal levels. Third, unlike Varela & Vermersch, we have no objection to abstract or imagined cases. As we see it, imagined cases (both abstract and detailed) have a role to play in personal-level investigations. Indeed, a central methodological tool of Husserlian phenomenology (which Varela & Vermersch invoke) is "imaginative variation": Phenomenology is the description not simply of actual experiences, but of possible, that is, imagined, ones. (This is what makes phenomenology a conceptual project, not simply a descriptive one.) Dennett's Marilyn example can be taken in precisely this spirit; hence their criticism of it is unfair. The example is designed to make a conceptual point, one that is unaffected by the variations in the perceptual situation they describe. Finally, although we agree that phenomenological accounts are needed to complement the natural scientific study of perception, and that they cannot be reduced to "heterophenomenology" (see Note 5), we see no reason to believe that, in the context of visual science, they need to include every affective detail about how the subject feels in perceiving the world.

Lomas argues against our statement that visual perception at the personal level is constrained by norms of rationality. He cites the Müller-Lyer illusion, which he thinks is a counterexample, because it persists despite the person's belief that it is an illusion. Our point, however, was not that having a perceptual experience is equivalent to forming a belief – the so-called belief-theory of perception (Armstrong 1968) – but rather that perceptual experience has a

rational bearing on judgment. How things look to someone is always *relevant* to how one *ought* to judge them as being; and this remains true even in cases where an illusion persists despite one's knowledge that it is an illusion. One cannot have the experience of the lines looking unequal without at least recognizing that the experience raises the question of whether the corresponding judgment – that the lines are unequal – is true. To say that perceptual experience and judgment can diverge and even contradict one another is to say that they can also be in accord, and that is enough to show the two kinds of content are normatively linked (see Noë 1995, where this view is first presented and elaborated).

NOTES
1. Author to whom all correspondence should be addressed.
2. We think this is the best interpretation of the passage from Teller and Pugh (1983, p. 581), which we quote in the target article. We would like to note, however, that in personal communication Teller has expressed doubts about this way of stating the bridge locus concept.
3. Although "locus" can mean any figure made up wholly of points that satisfy a given condition, rather than a place or locality, we believe that most cognitive neuroscientists conceptualize the bridge locus in the latter sense of a place or locality in the brain that is the immediate substrate of perceptual experience. Consider, for example, Crick and Koch's (1995) well-known view that the neural correlates of visual awareness have to be located somewhere in groups of cortical neurons that project directly to the frontal cortex, but not in V1.
4. Pishva claims that the homomorphism concept provides an alternative explanation for the Cumming & Friend (1980) study of the blind spot we discussed in section 7.1. He attributes to us the view that the mechanisms involved in blind spot completion may be at a higher level than those involved in the tilt aftereffect being probed. We did not endorse this view, however; we simply said it was not ruled out by their study. We fail to see the light that the homomorphism concept sheds on these experiments.
5. The requirement for the integration of contrast information has been discussed since the classic work of Land and McCann (1971) and Horn (1974). Recently, Ross and Pessoa (1997) have argued that previous brightness filling-in models (e.g., Grossberg & Todorović 1988) need to be extended to handle integration fully.
6. We wish to thank William Ross for conversations that have shaped the present discussion.
7. For **Dennett,** the person or animal would seem to be no more than a logical construct of subpersonal brain states and processes, just as the British Empire at the time of the War of 1812 was a logical construct of the King, the Members of Parliament, various officials and subjects of the Crown, and so on:

> Since You are nothing over and above the various subagencies and processes in your nervous system that compose you, the following sort of question is always a trap: "Exactly when did *I* (as opposed to various parts of my brain) become informed (aware, conscious) of some event?" Conscious experience, in our view, is a succession of states *constituted by* various processes occurring in the brain and not something over and above these processes that is *caused by* them (Dennett & Kinsbourne 1992, p. 236, emphasis in original).

Our view, however, is that the person or animal, although not something over and above its embodied being in the world (like a Cartesian Self), cannot be reduced to the various subagencies and processes of its nervous system (see Chiel & Beer 1997). The person or animal also has a different kind of intelligibility from its nervous system (requires different principles and descriptive vocabularies for its characterization), as Dennett (1987) himself has discussed, but seems unwilling to acknowledge fully (see Sedivy 1995).

8. Van Gulick thinks the issue here is about the philosophical doctrine of "phenomenal realism," which holds that "phenomenal mental states use a distinctive medium of representation, one that is often thought of as a plenum, a determinately differentiated continuous manifold." As he sees it, "[t]he intent of the Marilyns case is to call that view into question; there can be no such phenomenal representation because there is no underlying subpersonal representation with the detail to support it." Van Gulick is certainly right that **Dennett** criticizes realism about phenomenal consciousness. Our impression, however, is that Dennett rejects this doctrine because he thinks there is something "fishy," ill-defined, or incoherent about the very concept of phenomenal consciousness. If Van Gulick's interpretation is right, then Dennett would seem to be committed to the coherence of the notion of phenomenal consciousness after all, otherwise there would be no point to the argument that we lack the subpersonal resources needed to support that kind of consciousness.

9. In our view, such an account would not be merely "heterophenomenological," that is, it would not take as its data merely "the manifolds of subjects' [verbally or behaviorally expressed] beliefs about their conscious states" as interpreted from the perspective of the scientific observer; it would be phenomenological in the sense of also including as its data features of the experiences themselves (for further discussion see Pettitot et al. 1998 and Thompson et al. 1998).

10. As **Varela & Vermersch** note, such intersubjective descriptions are not available in the case of animal studies. This fact points toward an important difference between the *animal* and the *person*, a difference we have had to gloss over in our discussion of the personal level.

References

Letters "a" and "r" appearing before author's initials refer to target article and response, respectively.

Abeles, M. (1984) *Local circuits.* Springer Verlag. [aLP]
Adelson, E. H. (1993) Perceptual organization and the judgement of brightness. *Science* 262:2042–44. [CY]
Akins, K. & Winger, S. (1996) Ships in the night: Churchland and Ramachandran on Dennett's theory of consciousness. In: *Perception,* ed. K. Akins. Oxford University Press. [aLP]
Albright, T. D. (1995) "My most true mind thus makes mine eye untrue." *Trends in Neuroscience* 18:331–35. [aLP]
Anderson, B. L. (1995) New phenomena in surface interpolation. ARVO Abstract. *Investigative Ophthalmology and Visual Science* 36:S367. [arLP]
Anderson, B. L. & Julesz, B. (1995) A theoretical analysis of illusory contour formation in stereopsis. *Psychological Review* 102:705–43. [CY]
Appelbaum, I. (1995) *Speech perception: A philosophical analysis.* Doctoral Dissertation, University of Chicago. [IA]
Armstrong, D. M. (1968) *A materialist theory of mind.* Routledge & Kegan Paul. [rLP]
Arrington, K. F. (1994) The temporal dynamics of brightness filling-in. *Vision Research* 34:3371–87. [aLP]
(1996) Directional filling-in. *Neural Computation* 8(2):300–18. [KFA, aLP]
Assad, J. A. & Maunsell, J. H. R. (1993) Neuronal correlates of inferred motion in primate posterior parietal cortex. *Nature* 373:518–21. [aLP]
Bachmann, T. (1997) Visibility of brief images: The dual-process approach. *Consciousness and Cognition* 6:491–518. [TB]
Ballard, D. H. (1991) Animate vision. *Artificial Intelligence* 48:57–86. [arLP]
(1996) On the function of visual representation. In: *Perception,* ed. K. Akins. Oxford University Press. [arLP]

Ballard, D. H., Hayhoe, M. M., Pook, P. K. & Rao, R. P. N. (1997) Deitic codes for the embodiment of cognition. *Behavioral and Brain Sciences* 20(4) 723–67. [aLP]

Barlow, H. B. (1972) Single units and sensation. *Perception* 1:371–94. [WDR]
  (1995) The neuron doctrine in perception. In: *The cognitive neurosciences,* ed. M. Gazzaniga. MIT Press. [HBB]
  (1996) Banishing the homunculus. In: *Perception as a Bayesian inference,* ed. D. Knill & W. Richards. Cambridge University Press. [HBB]

Barlow, H. B. & Hill, R. (1963) Evidence for a physiological explanation of the waterfall phenomenon and figural aftereffects. *Nature* 200:1345–47. [aLP]

Barlow, H. B., Levick, W. R. & Yoon, M. (1971) Responses to single quanta of light in retinal ganglion cells of the cat. *Vision Research* 11(suppl. 3):87–101. [HBB]

Barlow, H. B. & Tripathy, S. P. (1997) Correspondence noise and signal pooling as factors determining the detectability of coherent visual motion. *Journal of Neuroscience* 17:7954–66. [HBB]

Baumgartner, G. (1961) Die Reaktionen der Neurone des zentralen visuellen systems der Katze im simultanen Helligkeitskontrast. In: *Neurophysiologie und Psychophysik des Visuellen Systems,* ed. R. Jung & H. H. Kornhuber. Springer. [LS]
  (1990) Where do visual signals become a perception? In: *The principles of design and operation of the brain,* ed. J. Eccles & O. Creutzfeldt. Pontificiae Academiae Scientiarum Scripta Varia 78. [LS]

Békéséy, G. von. (1968) Brightness distribution across the Mach bands measured with flicker photometry, and the linearity of sensory nervous interaction. *Journal of the Optical Society of America* 58:1–8. [aLP]

Berkeley, M., Debruyn, B. & Orban, G. (1994) Illusory, motion, and luminance-defined contours interact in the human visual system. *Vision Research* 34:209–16. [aLP]

Blackmore, S. J., Brelstaff, G., Nelson, K. & Troscianko, T. (1995) Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception* 24:1075–81. [JKO, arLP]

Blake, R. & Fox, R. (1974) Adaptation to invisible gratings and the site of binocular rivalry suppression. *Nature* 249:488–90. [FHD]

Block, N. (1995) Mental paint and mental latex. In: *Philosophical issues 7,* ed. E. Villanueve. Ridgeview. [AB]

Boselie, F. (1988) Local versus global minima in visual pattern completion. *Perception and Psychophysics* 43:431–45. [aLP]
  (1994) Local and global factors in visual occlusion. *Perception* 23:517–28. [aLP]

Boselie, F. & Wouterloud, D. (1992) A critical discussion of Kellman and Shipley's (1991) theory of occlusion phenomena. *Psychological Research/Psychologische Forschung* 54:278–85. [CY]

Boynton, R. (1983) Mechanisms of chromatic discrimination. In: *Colour vision: Physiology and psychophysics,* ed. J. Mollon & L. Sharpe. Academic Press. [aLP]

Boynton, R., Hayhoe, M. & Macleod, R. (1977) The gap effect: Chromatic and achromatic visual discrimination as affected by field separation. *Optica Acta* 24:159–77. [aLP]

Bridgeman, B. (1983) Isomorphism is where you find it. *Behavioral and Brain Sciences* 6:658–59. [aLP]

Brillouin, L. (1956) *Science and information theory.* Academic Press. [HBB]

Brooks, B. & Jung, R. (1973) Neuronal physiology of the visual cortex. In: *Handbook of sensory physiology, vol. VII/3: Central processing of visual information. Part B,* ed. R. Jung. Springer-Verlag. [TB]

Brooks, R. (1991) Intelligence without representation. *Artificial Intelligence* 47:139–59. [aLP]

Browman, C. & Goldstein, L. (1996) Dynamics and articulatory phonology. In: *Mind as motion: Explorations in the dynamics of cognition,* ed. R. Port & T. van Gelder. MIT Press. [IA]

Brown, R. & Thurmond, J. (1993) Preattentive and cognitive effects on perceptual completion at the blind spot. *Perception and Psychophysics* 53:200–09. [aLP]

Bruno, N., Bertamini, M. & Domini, F. (1997) Amodal completion of partly occluded surfaces: Is there a "mosaic" stage? *Journal of Experimental Psychology: Human Perception and Performance* 23(5):1412–26. [aLP]

Buffart, H., Leewenberg, E. & Restle, F. (1981) Coding theory of visual pattern completion. *Journal of Experimental Psychology: Human Perception and Performance* 7:241–74. [aLP]
  (1983) Analysis of ambiguity in visual pattern completion. *Journal of Experimental Psychology: Human Perception and Performance* 9:980–1000. [aLP]

Burkhardt, D. (1966) Brightness and the increment threshold. *Journal of the Optical Society of America* 56:9789–881. [aLP]

Campion, J., Latto, R. & Smith, Y. M. (1983) Is blindsight an effect of scattered light, spared cortex, and near-threshold vision? *Behavioral and Brain Sciences* 6:423–48. [KK]

Carman, G. J. & Welch, L. (1992) Three-dimensional illusory contours and surfaces. *Nature* 360:585–87. [BD]

Cavanagh, P. (1991) What's up in top-down processing? In: *Representations of vision: Trends and tacit assumptions in vision research,* ed. A. Gorea. Cambridge University Press. [DL]
  (1995) Vision is getting easier every day. *Perception* 24:1227–32. [DL]

Chapman, B. & Stone, L. S. (1996) Turning a blind eye to cortical receptive fields. *Neuron* 16:9–12. [aLP]

Chelazze, L., Miller, E. K., Duncan, J. & Desimone, R. (1983) A neural basis for visual search in inferior temporal cortex. *Nature* 363:345–47. [aLP]

Chiel, H. J. & Beer, R. D. (1997) The brain has a body: Adaptive behavior emerges from interactions of nervous system, body, and environment. *Trends in Neuroscience* 20:553–57. [rLP]

Chino, Y. M., Kaas, J. H., Smith, E. L., III, Langston, A. L. & Cheng, H. (1992) Rapid reorganization of cortical maps in adult cats following restricted deafferentation in retina. *Vision Research* 32:789–96. [IM]

Chino, Y. M., Smith, E. L., III, Kaas, J. H., Sasaki, Y. & Cheng, H. (1995) Receptive-field properties of deafferentated visual cortical neurons after topographic map reorganization in adult cats. *Journal of Neuroscience* 15:2417–33. [IM]

Churchland, P. S. & Ramachandran, V. S. (1993) Filling in: Why Dennett is wrong. In: *Dennett and his critics,* ed. B. Dahlbom. Basil Blackwell. Also in: *Perception,* ed. K. Akins. Oxford University Press, 1996. [aLP]

Cicerone, C. M. & Hoffman, D. D. (1991) Dynamic neon colors: Perceptual evidence for parallel visual pathways. *University of California, Irvine, Mathematical Behavioral Sciences Memo* 91–22. [MS]

Cicerone, C. M., Hoffman, D. D., Gowdy, P. D. & Kim, J. S. (1995) The perception of color from motion. *Perception and Psychophysics* 57:761–77. [TFS, MS]

Clark, A. (1993) *Sensory qualities.* Oxford University Press. [EM]
  (1996) *Being there: Putting brain, body, and world together again.* MIT Press. [arLP]

Cohen, M. & Grossberg, S. (1984) Neural dynamics of brightness perception: Features, boundaries, diffusion, and resonance. *Perception and Psychophysics* 36:428–56. [aLP]

Coren, S., Porac, C. & Theodor, L. H. (1986) The effects of perceptual set on the shape and apparent depth of subjective contours. *Perception and Psychophysics* 39(5):327–33. [DL]

Cornsweet, T. N. (1970) *Visual perception.* Academic Press. [aLP]

Cornsweet, T. N. & Teller, D. (1965) Relation of increment thresholds to brightness and luminance. *Journal of the Optical Society of America* 55:1303–08. [BD, aLP]

Cowey, A. (1979) Cortical maps and visual perception (The Grindley Memorial Lecture). *Quarterly Journal of Experimental Psychology* 31:1–7. [rLP]

Craik, K. J. W. (1940) Visual adaptation. Unpublished doctoral thesis, Cambridge University. [aLP]

Crane, T. (1992) The nonconceptual content of experience. In: *The contents of perception: Essays on perception,* ed. T. Crane. Cambridge University Press. [DL]

Creutzfeldt, O. D. (1990) Brain, perception, and mind. In: *Visual perception: The neurophysiological foundations,* ed. L. Spillmann & J. S. Werneer. Academic Press. [aLP]

Crick, F. & Coch, C. (1995) Are we aware of neural activity in primary visual cortex? *Nature* 375:121–23. [rLP]

Cumming, G. & Friend, H. (1980) Perception at the blind spot and tilt aftereffect. *Perception* 9:233–38. [arLP, VP]

Cunningham, D. W., Shipley, T. F. & Kellman, P. J. (1998) The perception of surface qualities in dynamic scenes: Spatiotemporally defined boundaries and surfaces. *Perception* 27:403–15 [TFS]

Damasio, A. (1994) *Descartes' error.* Grosset/Putnam. [EM]

Darian-Smith, C. & Gilbert, C. D. (1995) Topographic reorganization in the striate cortex of the adult cat and monkey is cortically mediated. *Journal of Neuroscience* 15:1631–47. [IM]

Davidson, M. & Whiteside, J. (1971) Human brightness perception near sharp contours. *Journal of the Optical Society of America* 61:530–36. [aLP]

Davis, G. (1996) *Attention and modal versus amodal completion: A function for filling-in?* Ph. D. thesis, Cambridge University. [PW]

Davis, G. & Driver, J. (1994) Parallel detection of Kanizsa subjective figures in the human visual system. *Nature* 371:791–93. [GD, rLP]
  (1997) Spreading of visual attention across modally versus amodally completed surfaces. *Psychological Science* 8:275–81. [GD]
  (1998a) A functional role for filling-in in the control of visual attention. *Perception* 26:1397–1412. [GD]
  (1998b) Kanizsa subjective figures can act as occluding surfaces in preattentive human vision. *Journal of Experimental Psychology: Human Perception and Performance* 8:275–81. [GD]

DeAngelis, G. C., Anzai, A., Ohzawa, I. & Freeman, R. (1995a) Receptive field structure in the visual cortex: Does selective stimulation induce plasticity? *Proceedings of the National Academy of Sciences (USA)* 92:9682–86. [aLP]

DeAngelis, G. C., Ohzawa, I. & Freeman, R. D. (1995b) Receptive-field dynamics in the central visual pathways. *Trends in Neuroscience* 18:451–57. [aLP]

Dement, W. C. (1976) *Some must watch while some must sleep.* Norton.   [AR]

Dennett, D. C. (1969) *Content and consciousness.* Routledge & Kegan Paul.   [DCD]

  (1978) Toward a cognitive theory of consciousness. In: *Brainstorms,* ed. D. C. Dennett. MIT Press/Bradford Books.   [AB, aLP]

  (1987) *The intentional stance.* MIT Press.   [rLP]

  (1991) *Consciousness explained.* Little, Brown/Allen Lane.   [HBB, AB, GD, PDW, FHD, LK, KK, JKO, arLP, AR]

  (1992) "Filling in" versus finding out: A ubiquitous confusion in cognitive science. In: *Cognition: Conceptual and methodological issues,* ed. H. L. Pick, Jr., P. van den Broek & D. C. Knill. American Psychological Association.   [AB, BD, KK, SL, aLP, VP, WDR, DT, CWT, RVG]

  (1993) Back from the drawing board. In: *Dennett and his critics,* ed. B. Dahlbom. Basil Blackwell.   [aLP]

  (1996a) Seeing is believing – or is it? In: *Perception,* ed. K. Akins. Oxford University Press.   [aLP]

  (1996b) Bewusstein hat mehr mit Ruhm als mit Fernsehen zu tun [Consciousness: More like fame than television]. In: *Die Technik auf dem Weg zur Seele,* ed. C. Maar, E. Poppel & T. Christaller. Rowohlt.   [LWH, aLP]

Dennett, D. C. & Kinsbourne, M. (1992) Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences* 15:183–247.   [arLP]

Denton, E. J. & Pirenne, M. H. (1954) The absolute sensitivity and functional stability of the human eye. *Journal of Physiology, London* 123:417–42.   [HBB]

De Valois, R., Webster, M. & De Valois, K. (1986) Temporal properties of brightness and color induction. *Vision Research* 26:887–97.   [aLP]

De Weerd, P., Desimone, R. & Ungerleider, L. (1998) Perceptual filling-in, a parametric study. *Vision Research* 38:2721–934.   [rLP, PDW]

De Weerd, P., Gattass, R., Desimone, R. & Ungerleider, L. (1995) Responses of cells in monkey visual cortex during perceptual filling-in of an artificial scotoma. *Nature* 377:731–34.   [PDW, LWH, HK, arLP, KS]

DeYoe, E. A. & Van Essen, D. C. (1988) Concurrent processing in the primate visual cortex. In: *The cognitive neurosciences,* ed. M. Gazzaniga. MIT Press.   [aLP]

Diamond, A. L. (1953) Foveal simultaneous brightness contrast as a function of inducing- and test-field luminances. *Journal of Experimental Psychology* 45:304–14.   [BD]

Dittmers, F. (1920) Über die Abhängigkeit der Unterschiedsschwelle für Helligkeiten von der antagonistischen Induktion. *Zeitschrift für Sinnesphysiologie* 51:214–32.   [BD]

Dresp, B. (1991) Etude psychophysique des mécanismes et processus de structuration dans la perception des formes. Unpublished Doctoral Thesis, Université Paris V.   [BD]

  (1992) Local brightness mechanisms sketch out surfaces but do not fill them in: Psychophysical evidence in the Kanizsa square. *Perception and Psychophysics* 52:562–70.   [BD, arLP]

  (1997) On "illusory" contours and their functional significance. *Current Psychology of Cognition* 16:489–518.   [BD]

Dresp, B. & Bonnet, C. (1991) Psychophysical evidence for low-level processing of illusory contours and surfaces in the Kanizsa square. *Vision Research* 31:1813–17.   [BD]

  (1993) Psychophysical measures of illusory form perception: Further evidence for local mechanisms. *Vision Research* 33:759–66.   [BD, rLP]

  (1995) Subthreshold summation with illusory contours. *Vision Research* 35:1071–78.   [BD]

Dresp, B. & Grossberg, S. (1997) Contour integration across polarities and spatial gaps: From contrast filtering to bipole cooperation. *Vision Research* 37:913–24.   [BD]

Durgin, F. H. (1995) Texture density adaptation and the perceived numerosity and distribution of texture. *Journal of Experimental Psychology: Human Perception and Performance* 21:149–69.   [FHD]

Durgin, F. & Cole, R. (1997) Texture-density aftereffects to filled-in and suppressed portions of textures. *Investigative Ophthalmology and Visual Science* 38:S636.   [FHD, rLP]

Durgin, F., Tripathy, S. P. & Levi, D. M. (1995) On the filling in of the visual blind spot: Some rules of thumb. *Perception* 24:827–40.   [KK, aLP, RW]

Dyken, M. E., Lin-Dyken, D. C., Seaba, P. & Yamada, T. (1995) Violent sleep-related behavior leading to subdural hemorrhage. *Archives of Neurology* 52:318–21.   [AR]

Epstein, W. (1982) Percept-percept coupling. *Perception* 11:75–83.   [TFS]

Eskew, R. (1989) The gap effect revisited: Slow changes in chromatic sensitivity as affected by luminance and chromatic borders. *Vision Research* 29:717–29.   [aLP]

Eskew, R. & Boynton, R. (1987) Effects of field area configuration on chromatic and border discrimination. *Vision Research* 27:1835–44.   [aLP]

Farah, M. J. (1988) Is visual imagery really visual? Overlooked evidence from neuropsychology. *Psychological Review* 95:307–17.   [DCE]

  (1990) *Visual agnosia: Disorders of object recognition and what they tell us about normal vision.* MIT Press.   [CY]

Fiorani, M., Rosa, M. G. P., Gattass, R. & Rocha-Miranda, C. E. (1992) Dynamic surrounds of receptive fields in primate striate cortex: A physiological basis for perceptual completion. *Proceedings of the National Academy of Sciences (USA)* 89:8547–51.   [FHD, HK, IM, arLP, KS]

Fiorentini, A. (1972) Mach band phenomena. In: *Handbook of sensory physiology, vol. VII-4,* ed. D. Jameson & L. Hurvich. Springer Verlag.   [BD, aLP]

Fiorentini, A. & Zoli, M. T. (1966) Detection of a target superimposed to a step pattern of illumination. *Atti della fondazione Giorgio Ronchi* 21:338–56.   [BD]

Fishman, M. C. & Michael, C. R. (1973) Integration of auditory information in cat's visual cortex. *Vision Research* 13:1415.   [aLP]

Fodor, J. (1975) *The language of thought.* Harvard University Press.   [rLP]

Fry, G. (1948) Mechanisms subserving simultaneous contrast. *American Journal of Optometry and Archives of the American Academy of Optometry* 25:162–78.   [aLP]

Gallistel, C. (1990) *The organization of learning.* MIT Press.   [VP, TFS]

Gattass, R., Fiorani, M., Rosa, M. G. P., Pinon, M. C. G., Barbosa de Sousa, A. P. & Soares, J. G. M. (1992) Visual responses outside the classical receptive field in primate striate cortex: A possible correlate of perceptual completion. In: *The visual system from genesis to maturity,* ed. R. Lent. Birkhauser.   [FHD]

Geldard, F. & Sherrick, C. (1972) The cutaneous "rabbit": A perceptual illusion. *Science* 178:178–79.   [aLP]

Gellatly, A. R. H. (1980) Perception of an illusory triangle with masked inducing figure. *Perception* 9:599–602.   [TB]

Gerbino, W. & Salmaso, D. (1987) The effect of amodal completion on visual matching. *Acta Psychologica* 65:25–46.   [CY]

Gerrits, H. J. M., de Haan, B. & Vendrick, A. (1966) Experiments with retinal stabilized images. Relations between the observations and neural data. *Vision Research* 6:427–40.   [aLP]

Gerrits, H. J. M. & Timmerman, G. J. M. E. N. (1969) The filling-in process in patients with retinal scotomata. *Vision Research* 9:439–42.   [aLP]

Gerrits, H. J. M. & Vendrik, A. J. H. (1970) Simultaneous contrast, filling-in process and information processing in man's visual system. *Experimental Brain Research* 11:411–30.   [aLP]

Gibson, J. J. (1972) A theory of direct visual perception. In: *The psychology of knowing,* ed. J. R. Royce & W. W. Rozeboom. Gordon & Breach.   [aLP]

  (1979) *The ecological approach to visual perception.* Houghton Mifflin.   [AB, FHD, arLP]

Gilbert, C. (1992) Horizontal integration and cortical dynamics. *Neuron* 9:1–13.   [arLP]

Gilbert, C. D. & Wiesel, T. N. (1990) The influence of contextual stimuli on the orientation selectivity of cells in the primary visual cortex of the cat. *Vision Research* 30:1689–1701.   [BD]

  (1992) Receptive field dynamics in adult primary visual cortex. *Nature* 356:150–52.   [HK, IM, aLP]

Gilchrist, A. L. (1977) Perceived lightness depends on perceived spatial arrangement. *Science* 195:185–87.   [CY]

Gillam, B. & Borsting, E. (1988) The role of monocular regions in stereoscopic displays. *Perception* 17:603–08.   [aLP]

Gillam, B. & McGrath, D. (1979) Orientation relative to the retina determines perceptual organization. *Perception and Psychophysics* 26:177–81.   [CY]

Gregory, R. L. (1966) *Eye and brain.* Weidenfeld and Nicolson.   [Fifth edition 1998, Oxford University Press.]   [RLG]

  (1972) Cognitive contours. *Nature* 238:51–52.   [aLP]

  (1997) Knowledge in perception and illusion. *Philosophical Transactions of the Royal Society of London B* 352:1121–28.   [RLG]

Grimes, J. (1996) On the failure to detect changes in scenes across saccades. In: *Perception,* ed. K. Akins. Oxford University Press.   [aLP]

Grosof, D., Shapley, R. & Hawken, M. (1993) Macaque V1 neurons can signal "illusory" contours. *Nature* 365:550–52.   [aLP]

Grossberg, S. (1983) The quantized geometry of visual space: The coherent computation of depth, form, and lightness. *Behavioral and Brain Sciences* 6:625–92.   [aLP]

  (1984) Outline of a theory of brightness, color, and form perception. In: *Trends in mathematical psychology,* ed. E. Degreef & J. Van Buggenhaut. Elsevier/North-Holland.   [SG]

  (1987a) Cortical dynamics of three-dimensional form, color, and brightness perception: I. Monocular theory. *Perception and Psychophysics* 41:87–116.   [SG, aLP]

  (1987b) Cortical dimensions of three-dimensional form, color, and brightness perception: II. Binocular theory. *Perception and Psychophysics* 41:117–58.   [aLP]

  (1994) 3-D vision and figure-ground separation by visual cortex. *Perception and Psychophysics* 55:48–120.   [SG, aLP, CY]

  (1997) Cortical dynamics of 3-D figure-ground perception of 2-D pictures. *Psychological Review* 104:618–58.   [SG]

(1998a) How does the cerebral cortex work? Learning, attention, and grouping by the laminar circuits of visual cortex. Technical Report CAS/CNS-TR-97–023. *Spatial Vision,* in press. [SG]

(1998b) The complementary brain: A new paradigm for brain specialization and modularity. Technical Report CAS/CNS-TR-98–003. (submitted). [SG]

Grossberg, S. & McLoughlin, N. (1997) Cortical dynamics of 3-D surface perception: I. Binocular and half-occluded scenic images. *Neural Networks* 9:1583–1605. [aLP]

Grossberg, S. & Mingolla, E. (1985) Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review* 92:173–211. [SG, arLP, WDR, CY]

Grossberg, S., Mingolla, E. & Ross, W. (1997) Visual brain and visual perception. *Trends in Neurosciences* 20:106–11. [rLP, WDR]

Grossberg, S. & Todorović, D. (1988) Neural dynamics of 1-D and 2-D brightness perception: A unified model of classical and recent phenomena. *Perception and Psychophysics* 43:241–77. [SG, arLP, DT]

Grüsser, O.-J. (1956) Reaktionen einzelner corticaler und retinaler Neurone der Katze auf Flimmerlicht und ihre Bezeihungen zur subjektiven Sinnesphysiologie. *Medizinische Dissertation,* Freiburg i. Br. [LS]

Haenny, P. E., Maunsell, J. H. R. & Schiller, P. H. (1988) State dependent activity in monkey visual cortex. *Experimental Brain Research* 69:245–59. [aLP]

Hamada, J. (1984) A multistage model for border contrast. *Biological Cybernetics* 39:81–86. [aLP]

Hardage, L. & Tyler, C. W. (1995) Induced twinkle aftereffect as a probe of dynamic visual processing mechanisms. *Vision Research* 35:757–66. [FHD, aLP]

Hardin, C. L. & Maffi, L., eds. (1997) *Color categories in thought and language.* Cambridge University Press. [rLP]

Hartline, H. K. (1940) The receptive fields of optic nerve fibers. *American Journal of Physiology* 130:690–99. [rLP]

(1949) Inhibition of activity of visual receptors by illuminating nearby retinal elements in the Limulus eye. *Federation Proceedings* 8:69. [rLP]

(1942) The neural mechanisms of vision. In: *The Harvey Lectures* 37:39–68.. The Harvey Society of New York. [LK]

He, S., Cavanagh, P. & Intriligator, J. (1996) Attentional resolution and the locus of visual awareness. *Nature* 383:334–37. [TB]

Heinemann, E. (1972) Simultaneous brightness induction. In: *Handbook of sensory physiology, vol. VII-4,* ed. D. Jameson & L. Hurvich. Springer Verlag. [aLP]

Held, R. R. & Hein, A. (1963) Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology* 56:872–76. [LWH]

Helmholtz, H. von (1925) *Physiological optics, vol. III. The theory of the perceptions of vision* (trans. from 3rd German edition, 1910). Optical Society of America. [HBB]

Hering, E. (1872/1964) *Outline of a theory of the light sense,* trans. L. M. Hurvich & D. J. Jameson. Harvard University Press. [LK, aLP]

Hoffman, D. D. (1998) *Visual intelligence.* Norton. [MS]

Horn, B. K. P. (1974) Determining lightness from an image. *Computer Graphics and Image Processing* 3:277–99. [rLP, WDR]

Horn, G. & Hill, R. M. (1969) Modifications of the receptive field of cells in the visual cortex occurring spontaneously and associated with bodily tilt. *Nature* 221:185–87. [aLP]

Horton, J. C. (1984) Cythochrome oxidase patches: A new cytoarchitectonic feature of monkey visual cortex. *Philosophical Transactions of the Royal Society of London* 304:199–253. [rLP]

Hubel, D. & Wiesel, T. (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology* 160:106–54. [aLP]

(1968) Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology* 195:215–43. [aLP]

Hurvich, L. & Jameson, D. (1957) An opponent process theory of color vision. *Psychological Review* 64:384–404. [LK]

Intraub, H. (1997) The representation of visual scenes. *Trends in Cognitive Sciences* 1:217–21. [rLP]

Jackson, F. (1982) Epiphenomenal qualia. *Philosophical Quarterly* 32:127–36. [DCE]

Julesz, B. (1971) *Foundations of cyclopean perception.* University of Chicago Press. [aLP]

Jung, R. (1973) Visual perception and neurophysiology. In: *Handbook of sensory physiology, vol. VII/3A. Central processing of visual information,* ed. H. Autrum, R. Jung, W. Loewenstein, D. M. MacKay & H. L. Teuber. Springer-Verlag. [LS]

Kaas, J. H., Krubitzer, L. A., Chino, Y. M., Langston, A. L., Polley, E. H. & Blair, N. (1990) Reorganization of retinotopic cortical maps in adult mammals after lesions of the retina. *Science* 248:229–31. [IM]

Kanizsa, G. (1955) Margini quasi-percettivi in campi con stimulazione omogenea. *Rivista di Psicologia* 49:7–30. [aLP]

(1979) *Organization in vision: Essays in gestalt perception.* Praeger Press. [FHD, aLP, CY]

Kanizsa, G. & Gerbino, W. (1982) Amodal completion: Seeing or thinking? In: *Organization and representation in perception,* ed. J. Beck. Erlbaum. [aLP]

Kapadia, M. K., Ito, M., Gilbert, C. D. & Westheimer, G. (1995) Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in V1 of alert monkeys. *Neuron* 15:843–56. [BD]

Kaufman, L. (1974) *Sight and mind.* Oxford University Press. [LK]

Kawabata, N. (1982) Visual information processing at the blind spot. *Perceptual and Motor Skills* 55:95–104. [aLP]

(1984) Perception at the blind spot and similarity grouping. *Perception and Psychophysics* 36:151–58. [aLP]

(1990) Structural information processing in peripheral vision. *Perception* 19:631–36. [aLP]

Kellman, P. & Loukides, M. (1987) An object perception approach to static and kinetic subjective contours. In: *The perception of illusory contours,* ed. S. Petry & G. E. Meyer. Springer-Verlag. [TFS]

Kellman, P. & Shipley, T. (1991) A theory of visual interpolation in object perception. *Cognitive Psychology* 23:141–221. [arLP, CY]

Kellman, P. J., Yin, C. & Shipley, T. F. (1998) A common mechanism for illusory and occluded figure completion. *Journal of Experimental Psychology: Human Perception and Performance* 24:859–69. [CY]

Kelso, J. A. S. (1995) *Dynamic patterns: The self organization of brain and behavior.* MIT Press. [rLP]

Kerr, N. (1993) Dreams of the blind. In: *Encyclopedia of sleep and dreaming,* ed. M. A. Carskadon. Macmillan. [AR]

Kingdom, F. & Moulden, B. (1989) Border effects on brightness: A review of findings, models and issues. *Spatial Vision* 3:225–62. [aLP]

Kinney, J. (1967) Color induction using asynchronous flashes. *Vision Research* 7:299–318. [aLP]

Knau, H. & Spillman, L. (1997) Brightness fading during *Ganzfeld* adaptation. *Journal of the Optical Society of America A* 14:1213–22. [LS]

Knill, D. C. & Kersten, D. (1991) Apparent surface curvature affects lightness perception. *Nature* 351:228–30. [CY]

Knudsen, E. L., du Lac, S. & Esterly, S. D. (1987) Computational maps in the brain. *Annual Review of Neuroscience* 10:41–65. [rLP]

Koffka, K. (1935) *Principles of gestalt psychology.* Harcourt Brace. [BD, LK, aLP, CY]

Köhler, W. (1920) *Die physischen Gestalten in Ruhe und im stationären Zustand* [*The physical gestalts in rest and in stationary states*]. Veiweg. [aLP]

(1929) An old psuedoproblem. In: *The selected papers of Wolfgang Köhler,* ed. M. Henle. Liveright. [Reprinted edition, 1971]. [aLP, AR]

(1930) The new psychology and physics. In: *The selected papers of Wolfgang Köhler,* ed. M. Henle. Liveright. [Reprinted edition, 1971]. [aLP]

(1947) *Gestalt psychology: An introduction to new concepts in modern psychology.* Liveright. [LK, SL, arLP]

(1960) The mind-body problem. In: *The selected papers of Wolfgang Köhler,* ed. M. Henle. Liveright. [Reprinted edition, 1971]. [aLP]

(1969) *The task of gestalt psychology.* Princeton University Press. [aLP]

Köhler, W. & Wallach, H. (1944) Figure after-effect, an investigation of visual processes. *Proceedings of the American Philosophical Society* 88:269–357. [LK]

Komatsu, H. & Kinoshita, M. (1997) The coding of figure size in the macaque striate cortex (V1). *Society for Neuroscience Abstracts* 23:454. [HK]

Komatsu, H. & Murakami, I. (1994) Behavioral evidence of filling-in at the blind spot of the monkey. *Visual Neuroscience* 11:1103–13. [aLP]

Komatsu, H., Murakami, I. & Kinoshita, M. (1995) Spatial summation properties of the macaque V1 neurons in the retinotopic representation of the blind spot. *Society for Neuroscience Abstracts* 21:1648. [HK]

(1996) Surface representation in the visual system. *Cognitive Brain Research* 5:97–104. [HK, IM, rLP]

Krauskopf, J. (1963) Effect of retinal image stabilization on the appearance of heterochromatic targets. *Journal of the Optical Society of America* 53:741–44. [aLP]

Kubovy, M. & Pomerantz, J. (1981) Perceptual organization: An overview. In: *Perceptual organization,* ed. M. Kubovy & J. Pomerantz. Erlbaum. [aLP]

Land, E. (1977) The retinex theory of color vision. *Scientific American* 237:108–28. [aLP]

Land, E. & McCann, J. (1971) Lightness and retinex theory. *Journal of the American Optical Society of America* 61:1–11. [rLP]

Lashley, K. S., Chow, K. L. & Semmes, J. (1951) An examination of the electric field theory of cortical integration. *Psychological Review* 58:123–36. [aLP]

Leopold, D. A. & Logothetis, N. K. (1996) Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature* 379:549–53. [TB]

Lesher, G. W. (1995) Illusory contours: Toward a neurally based perceptual theory. *Psychonomic Bulletin and Review* 2:279–321. [arLP]

LeVay, S., Connolly, M., Houde, J. & Van Essen, D. C. (1985) The complete patterns of ocular dominance stripes in the striate cortex and visual field of the macaque monkey. *Journal of Neuroscience* 5:486–501.   [rLP]

Liberman, A. M. & Mattingly, I. G. (1986) The motor theory of speech perception revised. *Cognition* 21:1–36.   [IA]

Lindberg, D. (1976) *Theories of vision from Al-Kindi to Kepler.* University of Chicago Press.   [EM]

Lockhead, G. R., Johnson, R. C. & Gold, F. M. (1980) Saltation through the blind spot. *Perception and Psychophysics* 27:545–49.   [aLP]

Logothetis, N. K., Leopold, D. A. & Sheinberg, D. L. (1996) What is rivalling during binocular rivalry? *Nature* 380:621–24.   [TB]

Mach, E. (1822) *Die Analyse der Empfindungen.* Fischer.   [BD]

(1865) Über die Wirkung der räumlichen Vertheilung des Lichterizes auf die Netzhaut, I. *Sitzungsberichte der mathematisch-naturwissenschaftlichen Classe der kaiserlichen Akademie der Wissenschaften* 52:303–22. Reprinted in: *Mach bands: Quantitative studies on neural networks in the retina,* ed. F. Ratliff. Holden Day, 1965.   [aLP]

MacKay, D. M. (1986) Vision – the capture of optical covariation. In: *Visual neuroscience,* ed. J. D. Pettigrew, K. J. Sanderson & W. R. Levick. Cambridge University Press.   [JKO]

Marr, D. (1982) *Vision.* W. H. Freeman.   [GD, FHD, DCE, arLP, WDR]

Martin, K. A. C. (1988) From enzymes to visual perception: A bridge too far? *Trends in Neurosciences* 9:380–84.   [aLP]

Mather, G. (1988) Temporal properties of apparent motion in subjective figures. *Perception* 17:729–36.   [BD]

Mattingley, J. B., Davis, G. & Driver, J. (1997) Visual completion in parietal extinction. *Science* 275:671–73.   [GD, rLP]

McClamrock, R. (1995) *Existential cognition. Computational minds in the world.* University of Chicago Press.   [aLP]

McClelland, J. L. & Elman, J. L. (1986) The TRACE model of speech perception. *Cognitive Psychology* 18:1–86.   [IA]

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

Metzger, W. (1935) *Gesetz des Schens.* [1953, 2nd edition.] Kramer.   [BD, LS]

Michotte, A., Thines, G. & Crabbe, G. (1964) *Les complements amodaux des structures perceptives. Studia Psycologica.* Publications Universitaires de Louvain.   [aLP, TFS]

Milner, A. D. & Goodale, M. A. (1995) *The visual brain in action.* Oxford University Press.   [FHD, rLP]

Mitchison, G. (1995) A type of duality between self-organizing maps and minimal wiring. *Neural Computation* 7:25–35.   [rLP]

Moran, J. & Desimone, R. (1985) Selective attention gates visual processing in extrastriate cortex. *Science* 229:782–84.   [aLP]

Moravec, L. & Beck, J. (1986) Amodal completion: Simplicity is not the explanation. *Bulletin of the Psychonomic Society* 24:269–72.   [aLP]

Morell, F. (1972) Visual system's view of acoustic space. *Nature* 238:44–46.   [aLP]

Movshon, J., Chambers, B. & Blakemore, C. (1972) Interocular transfer in normal humans, and those who lack stereopsis. *Perception* 1:483–90.   [aLP]

Müller, G. E. (1896) Zur Psychophysik der Gesichtsempfindungen [Conceiving the psychophysics of visual sensations] I. *Zeitschrift für Psychologie* 10:1–82.   [DCD, aLP]

Murakami, I. (1995) Motion aftereffect after monocular adaptation to filled-in motion at the blind spot. *Vision Research* 35:1041–45.   [FHD, IM, arLP]

Murakami, I., Komatsu, H. & Kinoshita, M. (1997) Perceptual filling-in at the artificial scotoma following a monocular retinal lesion in the monkey. *Visual Neuroscience* 14:89–101.   [HK, IM, rLP]

Mustillo, P. & Fox, R. (1986) The perception of illusory contours in the hypercyclopean domain. *Perception and Psychophysics* 40:362–63.   [aLP]

Myin, E. (1998) Holism, functionalism, and visual awareness. *Communication and Cognition* 31(1):3–20.   [EM]

Nakayama, K. (1994) Gibson: An appreciation. *Psychological Review* 101:353–56.   [rLP]

Nakayama, K. & Shimojo, S. (1990a) Toward a neural understanding of surface representation. *Cold Spring Harbor Symposia on Quantitative Biology* LV:911–24.   [arLP]

(1990b) Da Vinci stereopsis: Depth and subjective occluding contours from unpaired image points. *Vision Research* 30:1811–25.   [aLP]

(1992) Experiencing and perceiving visual surfaces. *Science* 257:1357–63.   [BD, HN]

Nakayama, K., Shimojo, S. & Ramachandran, V. S. (1990) Transparency: Relation to depth, subjective contours, luminance, and neon color spreading. *Perception* 19:497–513.   [CY]

Nakayama, K., Shimojo, S. & Silverman, G. H. (1989) Stereoscopic depth, its relation to image segmentation, grouping, and the recognition of occluded objects. *Perception* 18:55–68.   [GD]

Nerger, J., Piantanida, T. & Larimer, J. (1993) Color appearance of filled-in backgrounds affects hue cancellation, but not detection thresholds. *Vision Research* 33:165–72.   [aLP]

Neumann, H. (1996) Mechanisms of neural architecture for visual contrast and brightness perception. *Neural Networks* 9:921–36.   [aLP]

Neumann, H. & Pessoa, L. (1998a) Visual filling-in and surface property reconstruction. Submitted.   [rLP]

(1998b) Variations on filling-in. *Proceedings of the 1st Tübinger Wahrnehmungskonferenz, Tübingen, Germany, February 27-March 1.*   [HN]

Neumann, H., Sepp, W. & Mossner, P. (1997) Adaptive resonance in V1-V2 interaction: Grouping, illusory contours, and RF-organization. In: *Computational neuroscience – Trends in research 1997,* ed. J. Bower. Plenum Press.   [HN, rLP]

Newsome, W. T., Britten, K. H., Salzman, C. D. & Movshon, J. A. (1990) Neuronal mechanisms of motion perception. *Cold Spring Harbor Symposia on Quantitative Biology* 55:697–705.   [HBB]

Noë, R. A. (1995) *Experience and the mind: An essay on the metaphysics of perception.* Doctoral dissertation, Department of Philosophy, Harvard University.   [arLP]

Noë, R. A., Pessoa, L. & Thompson, E. (forthcoming) Beyond the grand illusion: What change blindness really teaches us about vision.   [rLP]

Nusbaum, H. C. & Henley, A. S. (in press) Understanding speech perception from the perspective of cognitive psychology. In: *Theories in spoken language: Perception, production, and development,* ed. J. Charles-Luce, P. A. Luce & J. R. Sawusch. Ablex Publishing.   [IA]

O'Brien, V. (1958) Contour perception, illusion, and reality. *Journal of the Optical Society of America* 48:112–19.   [aLP]

O'Regan, K. (1992) Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology* 46:461–88.   [SL, JKO, arLP]

O'Regan, J. K., Rensink, R. A. & Clark, J. J. (1996) "Mud splashes" render picture changes invisible. ARVO Abstract. *Investigative Ophthalmology and Visual Science* 37:S213.   [JKO, arLP]

Palmer, S. & Rock, I. (1994) Rethinking perceptual organization: The role of uniform connectedness. *Psychonomic Bulletin and Review* 1:29–55.   [CY]

Paradiso, M. A. & Nakayama, K. (1991) Brightness perception and filling-in. *Vision Research* 31:1221–36.   [TB, PDW, aLP, KS]

Paradiso, M. A., Shimojo, S. & Nakayama, K. (1989) Subjective contours, tilt-aftereffects, and visual cortical organization. *Vision Research* 29:1205–13.   [aLP]

Parker, A. & Hawken, M. (1985) The capabilities of monkey cortical cells in spatial resolution tasks. *Journal of the Optical Society of America* A 2:1101–14.   [HBB]

Pessoa, L. (1996a) Mach bands: How many models are possible? Recent experimental findings and modeling attempts. *Vision Research* 36:3205–27.   [arLP]

(1996b) Mach band attenuation by adjacent stimuli: Experiment and filling-in simulations. *Perception* 25:425–42.   [rLP]

Pessoa, L., Mingolla, E. & Neumann, H. (1995) A contrast- and luminance-driven multiscale network model of brightness perception. *Vision Research* 35:2201–23.   [arLP]

Pessoa, L. & Neumann, H. (1998) Why does the brain fill-in? *Trends in Cognitive Sciences* 2:422–24.   [rLP]

Peterhans, E. & von der Heydt, R. (1989) Mechanisms of contour perception in monkey visual cortex. II. Contours bridging gaps. *Journal of Neuroscience* 9:1749–63.   [aLP]

Peterhans, E., von der Heydt, R. & Baumgartner, G. (1986) Neuronal responses to illusory contour stimuli reveal stages of visual cortical processing. In: *Visual neuroscience,* ed. J. D. Pettigrew, K. J. Sanderson & W. R. Levick. Cambridge University Press.   [rLP]

Peterson, M. A. (1994) The proper placement of uniform connectedness. *Psychonomic Bulletin and Review* 1:509–14.   [CY]

Peterson, M. A. & Gibson, B. S. (1994) Object recognition contributions to figure-ground organization: Operations on outlines and subject contours. *Perception and Psychophysics* 56(5):551–64.   [DL]

Petitot, J., Roy, J-M., Pachoud, B. & Varela, F. J., eds. (1998) *Naturalizing phenomenology: Issues in contemporary phenomenology and cognitive science.* Stanford University Press.   [rLP]

Petitot, J., Varela, F., Pachoud, B. & Roy, M., eds. (1999) *Naturalizing phenomenology: Current issues in phenomenology and cognitive science.* Stanford University Press.   [FJV]

Petter, G. (1956) Nuove richerche sperimentali sulla totalizzazione percettiva. *Revista di Psicologia* 50:213–27.   [aLP]

Pettet, M. W. & Gilbert, C. D. (1992) Dynamic changes in receptive-field size in cat primary visual cortex. *Proceedings of the National Academy of Sciences (USA)* 89:8366–70.   [aLP, KS]

Peugeot, C. (1999) The intuitive experience. *Journal of Consciousness Studies* 5(4). (in press).   [FJV]

Piatandia, T. (1985) Temporal modulation sensitivity of the blue mechanism: Measurements made with extraretinal chromatic adaptation. *Vision Research* 25:1439–44.   [aLP]

Poggio, T., Torre, V. & Koch, C. (1985) Computational vision and regularization theory. *Nature* 317:314–19.   [arLP]

Pöppel, E. & Richards, W. (1974) Light sensitivity in cortical scotomata contralateral to small islands of blindness. *Experimental Brain Research* 21:125–30.   [RW]

Poppelreuter, W. (1917/1990) *Disturbances of lower and higher visual capacities caused by occipital damage: With special reference to the psychopathological, pedagogical, industrial, and social implications.* Clarendon Press.   [RW]

Port, R. F. & Van Gelder, T., eds. (1995) *Mind as motion: Explorations in the dynamics of cognition.* MIT Press.   [rLP]

Pradzny, K. (1985) On the nature of inducing forms generating perceptions of illusory contours. *Perception and Psychophysics* 37:175–78.   [aLP]

Pugh, E. & Mollon, J. (1979) A theory of the pi 1 and pi 3 color mechanisms of Stiles. *Vision Research* 19:779–88.   [aLP]

Pylyshyn, Z. W. (1981) The imagery debate: Analogue media versus tacit knowledge. *Knowledge* 88:16–45.   [DCE]

Ramachandran, V. S. (1992a) Blind spots. *Scientific American* 266:86–91.   [aLP]

   (1992b) Filling in gaps in perception: Part 1. *Current Directions in Psychological Science* 1:199–205.   [aLP]

   (1993a) Filling in gaps in perception: Part 2. *Current Directions in Psychological Science* 2:56–65.   [aLP]

   (1993b) Filling in gaps in logic: Some comments on Dennett. *Consciousness and Cognition* 2:165–68.   [aLP]

   (1993c) Behavioral and magnetoencephalographic correlates of plasticity in the adult human brain. *Proceedings of the National Academy of Sciences USA* 90:10413–20.   [HK]

   (1995) Filling in gaps in logic: Reply to Durgin et al. *Perception* 24:841–45.   [LWH, aLP]

Ramachandran, V. S. & Gregory, R. (1991) Perceptual filling-in of artificially induced scotomas in human vision. *Nature* 350:699–702.   [DCD, arLP, KS]

Ramachandran, V. S., Gregory, R. & Aiken, W. (1993) Perceptual fading of visual texture borders. *Vision Research* 33:717–22.   [arLP]

Ratliff, F. (1965) *The Mach bands: Quantitative studies on neural networks in the retina.* Holden Day.   [LK]

Ratliff, F. & Sirovich, L. (1978) Equivalence classes of visual stimuli. *Vision Research* 18:845–51.   [SG, aLP, WDR]

Rechtschaffen, A. & Buchignani, C. (1992) The visual appearance of dreams. In: *The neuropsychology of sleep and dreaming,* ed. J. S. Antrobus & M. Bertini. Erlbaum.   [AR]

Rechtschaffen, A. & Foulkes, D. (1965) Effect of visual stimuli on dream content. *Perceptual and Motor Skills* 20:1149–60.   [AR]

Redies, C., Crook, J. & Creutzfeld, O. (1986) Neuronal responses to borders with and without luminance gradients in cat visual cortex and dorsal lateral geniculate nucleus. *Experimental Brain Research* 61:469–81.   [aLP]

Redies, C. & Spillman, L. (1981) The neon color effect in the Ehrenstein illusion. *Perception* 10:667–81.   [aLP]

Reid, R. C. & Shapley, R. (1988) Brightness induction by local contrast and the spatial dependence of assimilation. *Vision Research* 28:115–32.   [BD]

Rensink, R. A. & Enns, J. T. (1998) Early completion of occluded objects. *Vision Research* 38:2489–2505.   [rLP, CY]

Rensink, R. A., O'Regan, J. K. & Clark, J. J. (1996) To see or not to see: The need for attention to perceive changes in scenes. ARVO Abstract. *Investigative Ophthalmology and Visual Science* 37:S213.   [aLP]

Revensuo, A. (1993) Is there a ghost in the cognitive machinery? *Philosophical Psychology* 6:387–405.   [AR]

   (1994) The "multiple drafts" model and the ontology of consciousness. *Behavioral and Brain Sciences* 17:177–78.   [AR]

   (1995) Consciousness, dreams, and virtual realities. *Philosophical Psychology* 8:35–58.   [AR]

   (1997) How to take consciousness seriously in cognitive neuroscience. *Communication and Cognition* 30:185–206.   [AR]

Ringach, D. L. & Shapley, R. (1996) Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Research* 36:3037–50.   [CY]

Rizzolatti, G., Riggio, L. & Sheliga, B. M. (1994) Space and selective attention. In: *Attention and performance XV. Conscious and nonconscious information processing,* ed. C. Umiltà & M. Moscovitsch. MIT Press.   [rLP]

Rock, I. & Anson, R. (1979) Illusory contours as a solution to a problem. *Perception* 8:665–81.   [aLP]

Ross, W. D. & Pessoa, L. (1997) The selective integration neural network model of lightness perception. *Proceedings of the International Conference on Neural Networks (ICNN'97), Houston, Texas, USA, June, 9–12.*   [arLP, WDR]

Rossi, A. & Paradiso, M. A. (1996) Temporal limits of brightness induction and mechanisms of brightness perception. *Vision Research* 36:1391–98.   [aLP]

Rossi, A. F., Rittenhouse, C. D. & Paradiso, M. A. (1996) The representation of brightness in primary visual cortex. *Science* 273:1104–07.   [HK, arLP, DT]

Saunders, B. A. C. & van Brakel, J. (1997) Are there nontrivial constraints on colour categorization? *Behavioral and Brain Sciences* 20:167–228.   [rLP]

Scheerer, E. (1994) Psychoneural isomorphism: Historical background and current relevance. *Philosophical Psychology* 7:183–210.   [aLP]

Schenk, C. H. (1993) REM sleep behavior disorder. In: *Encyclopedia of sleep and dreaming,* ed. M. A. Carskadon. Macmilllan.   [AR]

Schiller, P. (1995) Effects of lesions in visual cortical area V4 on the recognition of transformed objects. *Nature* 376:342–44.   [aLP]

Schwartz, E. (1980) Computational anatomy and functional architecture of striate cortex: A spatial mapping approach to perceptual coding. *Vision Research* 20:645–69.   [rLP]

Searle, J. R. (1997) *The mystery of consciousness.* New York Review.   [AR]

Sedivy, S. (1995) Consciousness explained: Ignoring Ryle & Co. *Canadian Journal of Philosophy* 25:455–83.   [arLP]

Sekuler, A. B. (1994) Local and global minima in visual completion: Effects of symmetry and orientation. *Perception* 23:529–45.   [aLP]

Sekuler, A. B. & Blake, R. (1985) *Perception.* Alfred A. Knopf.   [aLP]

Sekuler, A. B. & Palmer, S. E. (1992) Perception of partly occluded objects: A microgenetic analysis. *Journal of Experimental Psychology: General* 121:95–111.   [aLP, CY]

Shepard, R. N. (1981) Psychophysical complementarity. In: *Perceptual organization,* ed. M. Kubovy & J. R. Pomerantz. Erlbaum.   [CWT]

Sheperd, G. M. (1983, 1988, 1994) *Neurobiology,* Third Edition. Oxford University Press.   [aLP]

Sherrington, C. S. (1906) *The integrative action of the nervous system.* Yale University Press.   [PW]

Shipley, T. F. & Kellman, P. J. (1992) Perception of partly occluded objects and illusory figures: Evidence for an identity hypothesis. *Journal of Experimental Psychology: Human Perception and Performance* 18:106–20.   [GD, aLP]

   (1994) Spatiotemporal boundary formation: Boundary, form, and motion perception from transformation of surface elements. *Journal of Experimental Psychology: General* 123:3–20.   [TFS]

   (1997) Spatiotemporal boundary formation: The role of local motion signals in boundary perception. *Vision Research* 37:1281–93.   [TFS]

Sigman, E. & Rock, I. (1974) Stroboscopic movement based on perceptual intelligence. *Perception* 3:9–28.   [TFS]

Simons, J. S. & Levin, D. T. (1997) Change blindness. *Trends in Cognitive Sciences* 1:261–67.   [rLP]

Singer, W. (1995) Time as coding space in neocortical processing: A hypothesis. In: *The cognitive neurosciences,* ed. M. Gazzaniga. MIT Press.   [aLP]

Singh, M. & Hoffman, D. D. (1997) Constructing and representing visual objects. *Trends in Cognitive Science* 1:98–102.   [MS]

Skinner, B. F. (1963) Behaviorism at fifty. *Science* 140:951–58.   [DCE]

Smith, A. & Over, R. (1975) Tilt aftereffects with subjective contours. *Nature* 257:581–82.   [aLP]

   (1976) Color-selective tilt aftereffects with subjective contours. *Perception and Psychophysics* 20:305–08.   [aLP]

   (1977) Orientation masking and the tilt illusion with subjective contours. *Perception* 6:441–47.   [aLP]

   (1979) Motion aftereffect with subjective contours. *Perception and Psychophysics* 25:95–98.   [aLP]

Smythies, J. (1996) A note on the concept of the visual field in neurology, psychology, and visual neuroscience. *Perception* 25:369–71.   [AR]

Sperry, R. W. & Miner, N. (1955) Pattern perception following insertion of mica plates into the cortex. *Journal of Comparative and Physiological Psychology* 48:463–69.   [aLP]

Sperry, R. W., Miner, N. & Myers, R. E. (1955) Visual pattern perception following subplial slicing and tantalum wire implantations in the visual cortex. *Journal of Comparative and Physiological Psychology* 48:50–58.   [aLP]

Spillman, L. (1981) Illusions of contrast, brightness, color, and motion and their neurophysiological interpretation. *Freiburger Universitätsblätter* 74:73–77.   [BD]

Spillman, L. & Dresp, B. (1995) Phenomena of illusory form: Can we bridge the gap between levels of explanation? *Perception* 24:1333–64.   [aLP]

Spillman, L., Ransom-Hogg, A. & Oehler, R. (1987) A comparison of perceptive and receptive fields in man and monkey. *Human Neurobiology* 6:51–62.   [LS]

Spillman, L. & Werner, J. S. (1996) Long-range interactions in visual perception. *Trends in Neurosciences* 19:428–34.   [rLP, LS]

Srinivasan, M. V., Chahl, J. S., Nagle, M. G. & Zhang, S. W. (1997) Embodying natural vision into machines. In: *From living eyes to seeing machines,* ed. M. V. Srinivasan & S. Venkatesh. Oxford University Press.   [HN]

Steinman, R. M. & Levinson, J. Z. (1990) The role of eye movement in the detection of contrast and spatial detail. In: *Eye movements and their role in visual and cognitive processes,* ed. E. Kowler. Elsevier Science Publishers.   [LWH]

Stevens, K. N. & Blumstein, S. E. (1981) The search for invariant acoustic

correlates of phonetic features. In: *Perspectives on the study of speech,* ed. P. D. Eimas & J. L. Miller. Erlbaum.   [IA]

Stoper, A. E. & Mansfield, J. G. (1978) Metacontrast and paracontrast suppression of a contourless area. *Vision Research* 18:1669–74.   [TB, aLP]

Takeichi, H., Nakazawa, H., Murakami, L. & Shimojo, S. (1995) The theory of the curvature-constraint line for amodal completion. *Perception* 24:373–89.   [aLP, CY]

Takeichi, H., Shimojo, S. & Watanabe, T. (1992) Neon flank and illusory contour: Interaction between two processes leads to color filling-in. *Perception* 21:313–24.   [aLP]

Teller, D. Y. (1980) Locus questions in visual science. In: *Visual coding and adaptability,* ed. C. Harris. Erlbaum.   [aLP]
    (1984) Linking propositions. *Vision Research* 24:1233–46.   [aLP]
    (1998) Linking propositions in visual science. Presented at "Toward a Science of Consciousness 1998: Tuscon III," University of Arizona, Tuscon, April 27, 1998.   [rLP]
    (1990) The domain of visual science. In: *Visual perception: The neurophysiological foundations,* ed. L. Spillman & J. S. Werneer. Academic Press.   [aLP]

Teller, D. Y. & Pugh, E. N., Jr. (1983) Linking propositions in color vision. In: *Colour vision: Physiology and psychophysics,* ed. J. D. Mollon & L. T. Sharpe. Academic Press.   [HBB, LWH, arLP, WDR, CWT, RVG]

Thompson, E. (1995) *Colour vision: A study in cognitive science and the philosophy of perception.* Routledge Press.   [AB, arLP]

Thompson, E., Noë, A. & Pessoa, L. (1999) Perceptual completion: A case study in phenomenology and cognitive science. In: *Naturalizing phenomenology: Issues in contemporary phenomenology and cognitive science,* ed. J. Petitot, J-M. Roy, B. Pachoud & F. J. Varela. Stanford University Press.   [arLP, FJV]

Thompson, E., Palacios, A. & Varela, F. J. (1992) Ways of coloring: Comparative color vision as a case study for cognitive science. *Behavioral and Brain Sciences* 15:1–74.   [aLP]

Todorović, D. (1987) The Craik-O'Brien-Cornsweet effect: New varieties and theoretical implications. *Perception and Psychophysics* 42:545–600.   [aLP, VP, DT]

Torjussen, T. (1978) Visual processing in cortically blind hemifield. *Neuropsychologia* 16:15–21.   [RW]

Treue, S. & Maunsell, J. H. R. (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382:539–41.   [aLP]

Trick, L. M. & Enns, J. T. (1997) Clusters precede shapes in perceptual organization. *Psychological Science* 8:124–29.   [CY]

Tripathy, S. & Levi, D. (1994) Long-range dichoptic interaction in the human visual cortex in the region corresponding to the blind spot. *Vision Research* 34:1127–38.   [arLP]

Tripathy, S., Levi, D., Ogmen, H. & Harden, C. (1995) Perceived length across the physiological blind spot. *Visual Neuroscience* 12:385–402.   [arLP]

Turvey, M. T., Shaw, R. E., Reed, E. S. & Mace, W. M. (1981) Ecological laws of perceiving and acting: In reply to Fodor and Pylyshyn (1981). *Cognition* 9:237–304.   [aLP]

Vaadia, E., Haalman, I., Abeles, M., Bergman, H., Prut, Y., Slovin, H. & Aertsen, A. (1995) Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* 373:515–18.   [aLP]

Vallortigara, G. (1987) The hidden face of the Kanizsa triangle: Apparent movement of subjective figures in three-dimensional space. *Perception* 16:449–52.   [BD]

Van Esen, J. S. & Novak, S. (1974) Detection thresholds within a display that manifests contour enhancement and brightness contrast. *Journal of the Optical Society of America* 64:726–29.   [BD]

van Tuijl, H. F. J. M. (1975) A new visual illusion: Neonlike color spreading and complementary color induction between subjective contours. *Acta Psychologica* 39:441–45.   [aLP]

Varela, F. J. (1995) Resonant cell assemblies. *Biological Research* 28:81–95.   [aLP]
    (1996) Neurophenomenology: A methodological remedy to the hard problem. *Journal of Consciousness Studies* 3:330–50.   [FJV]

Varela, F. J. & Bourgine, P. (1992) *Toward a practice of autonomous systems. Proceedings of the First European Conference on Artificial Life.* MIT Press.   [aLP]

Varela, F. J. & Shear, J., eds. (1999) The view from within: First-person methodologies for the study of consciousness: Special issue. *Journal of Consciousness Studies* 5(4). (in press).   [FJV]

Varela, F. J., Thompson, E. & Rosch, E. (1991) *The embodied mind: Cognitive science and human experience.* MIT Press.   [aLP]

Vermersch, P. (1998) Introspection as practice. *Journal of Consciousness Studies* 5(4). (in press).   [FJV]

von der Heydt, R. & Peterhans, E. (1989) Mechanisms of contour perception in monkey visual cortex. I. Lines of pattern discontinuity. *Journal of Neuroscience* 9:1731–48.   [aLP]

von der Heydt, R., Peterhans, E. & Baumgartner, G. (1984) Illusory contours and cortical neuron responses. *Science* 224:1260–62.   [aLP]

von der Heydt, R., Zhou, H. & Friedman, H. S. (1996) The coding of extended colored figures in monkey visual cortex. *Society for Neuroscience Abstract* 22:951.   [HK, rLP]

von Grünau, M. (1979) The involvement of illusory contours in stroboscopic motion. *Perception and Psychophysics* 25:205–08.   [BD]

Walker, R. & Mattingley, J. B. (1997) Ghosts in the machine? Pathological visual completion phenomena in the damaged brain. *Neurocase* 3:313–35.   [rLP, RW]

Walker, R. & Young, A. W. (1996) Object-based neglect: An investigation of the contributions of eye movements and perceptual completion. *Cortex* 32:279–95.   [RW]

Wallach, H. (1935/1996) "On the visually perceived direction of motion" by Hans Wallach: 60 years later. (Trans. S. Wuerger, R. Shapley & N. Rubin). *Perception* 25:1317–67.   [TFS]

Walls, G. (1954) The filling-in process. *American Journal of Optometry* 31:329–40.   [aLP, PW]

Warrington, E. K. (1962) The completion of visual forms across hemianopic field defects. *Journal of Neurology, Neurosurgery and Psychiatry* 25:208–17.   [RW]

Watanabe, T. & Cavanagh, P. (1991) Texture and motion spreading, the aperture problem, and transparency. *Perception and Psychophysics* 50(5):459–64.   [arLP, KS]
    (1993) Surface decomposition accompanying the perception of transparency. *Spatial Vision* 7:95–111.   [arLP]

Watanabe, T. & Sato, T. (1989) Effects of luminance contrast on color spreading and illusory contour in the neon color spreading effect. *Perception and Psychophysics* 45:427–30.   [aLP]

Wehrhahn, C. & Dresp, B. (1998) Detection facilitation by collinear stimuli in humans: Dependence on strength and sign of contrast. *Vision Research* 38:423–28.   [BD]

Weiskrantz, L. (1990) The Ferrier Lecture 1989. Outlooks for blindsight: Explicit methodologies for implicit process. *Proceedings of the Royal Society of London B* 239:247–78.   [RW]
    (1996) Blindsight revisited. *Current Opinions in Neurobiology* 6:215–20.   [KK]

Weisstein, N. (1969) What the frog's eye tells the human brain: Single cell analyzers in the human visual system. *Psychological Bulletin* 72:157–76.   [aLP]

Wertheimer, M. (1912) Experimentelle Studien über das Sehen von Bewegung [Experimental studies on the visual perception of motion]. *Zeitschrift für Psychologie* 61:161–265.   [aLP, LS]

White, K. D., Petry, H., Riggs, L. A. & Miller, J. (1978) Binocular interactions during establishment of McCollough effects. *Vision Research* 18:1201–15.   [FHD]

Wildman, K. N. (1974) Visual sensitivity at an edge. *Vision Research* 14:749–53.   [BD]

Williams, L. R. & Hanson, A. R. (1996) Perceptual completion of occluded surfaces. *Computer Vision and Image Understanding* 64:1–20.   [CY]

Wouterloud, D. & Boselie, F. (1992) A good-continuation model of some occlusion phenomena. *Psychological Research* 54:267–77.   [aLP]

Wu, M. F. (1993) Sensory processing and sensation during sleep. In: *Encyclopedia of sleep and dreaming,* ed. M. A. Carskadon. Macmillan.   [AR]

Yarbus, A. L. (1967) *Eye movements and vision.* Plenum Press.   [aLP]

Yin, C., Kellman, P. J. & Shipley, T. F. (1997) Surface completion complements boundary interpolation. *Perception* 26:1459–79.   [Special issue: *Contextual effects on color appearance II*].   [CY]

Yu, C. & Levi, D. M. (1997) Spatial facilitation predicted with end-stopped spatial filters. *Vision Research* 37:3117–27.   [BD]

Zeki, S. (1983a) Colour coding in the cerebral cortex: The reactions of cells in monkey visual cortex to wavelengths and colours. *Neuroscience* 9:741–65.   [aLP]
    (1983b) Colour coding in the cerebral cortex: The responses of wavelength-selective cells in monkey visual cortex to changes in wavelength composition. *Neuroscience* 9:767–81.   [aLP]

Zeki, S. & Shipp, S. (1988) The functional logic of cortical connections. *Nature* 335:311–17.   [aLP]