12

Filling In: Why Dennett Is Wrong

Patricia S. Churchland and V. S. Ramachandran

It comes as a surprise to discover that the foveal area in which one has high resolution and high acuity vision is minute: it encompasses a mere 2 degrees of visual angle—roughly, the area of a thumbnail at arm's length. The introspective guess concerning acuity in depth likewise errs on the side of extravagance; the region of crisp, fused perception is, at arm's length, only a few centimeters deep; closer in, the area of fused perception is even narrower. The eyes make a small movement—a saccade—about every 200 to 300 milliseconds, sampling the scene by continually shifting the location of the fovea. Presumably, interpolation across intervals of time to yield an integrated spatiotemporal representation is a major component of what brains do. Interpolation in perception probably enables generation of an internal representation of the world that is useful in the animal's struggle for survival.

The debut demonstration of the blind spot in the visual field is comparably surprising. The standard setup requires monocular viewing of an object offset about 13 to 15 degrees from the point of fixation (figure 12.1). If the object falls in the region of the blind spot of the viewing eye, the object will not be perceived. Instead, the background color and texture will be seen as uniform across the region. This is generally characterized as "filling in" of the blind spot. The existence of the perceptual blind spot is owed to the specific architecture of the retina. As shown in figure 12.2, each retina has a region where the optic nerve leaves the retina and hence where no transducers (rods and cones) exist. This region is the blind spot. Larger than the fovea, it is about 6.0 degrees in length and about 4.5 degrees in width.

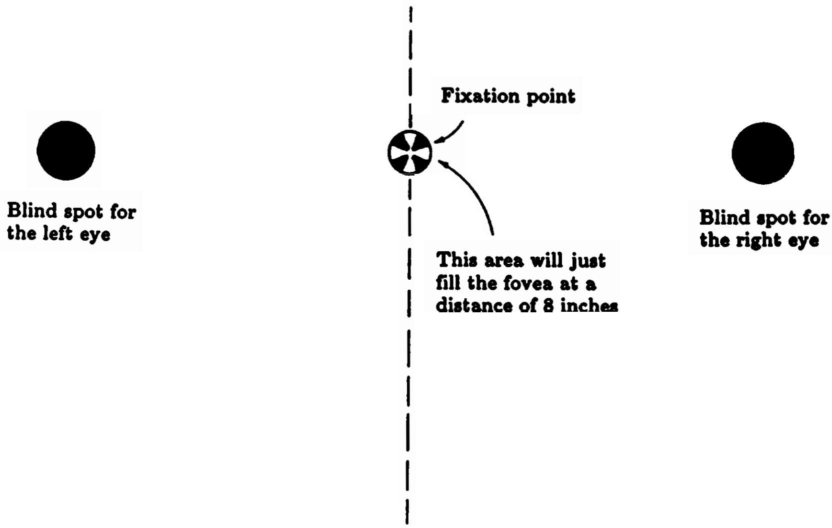

Figure 12.1
Instructions: close your right eye. Hold the page about eight inches in front of you. Hold the page very straight, without tilting or rotating it. Stare at the fixation point. Adjust the distance of the page until the black spot on the left disappears. Repeat with the left eye closed. (After Lindsay and Norman, 1972.)

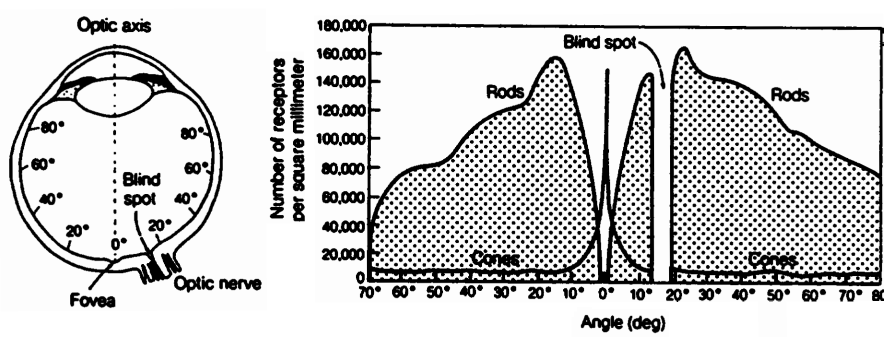

Figure 12.2
(*Left*) The human eye. The blind spot (optic disk) is that region on the retina where the ganglion cells leave the retina and project to the lateral geniculate nucleus of the thalamus. (*Right*) The packing density of light-sensitive cones is greatest at the fovea, decreasing sharply in the peripheral field. Rod density is greatest in the region that immediately surrounds the fovea, and gradually decreases in the more peripheral regions. Note that the region of the blind spot is in the peripheral field, and is larger than the foveal area.

Relying on two eyes, a perceiver—even a careful and discerning perceiver—will fail to notice the blind spot, mainly because the blind regions of the two eyes do not overlap. If light from a thimble, for example, falls in the blind spot of the left eye, it will nevertheless be detected normally by the right retina, and the viewer sees a thimble. Even in the monocular condition, however, one may fail to notice the blind spot because objects whose borders extend past the boundaries of the blind spot tend to be seen as filled in, as without gaps.

I. Dennett's Hypothesis

What is going on when one's blind spot is seen as filled in—as without gaps in the scene? Is it analogous to acquiring the nonvisual representation (belief) that Bowser, the family dog, is under the bed, on the basis of one's visual perception of his tail sticking out? Or is it more akin to regular visual perception of the whole Bowser in one's *peripheral but nonblind field*? That is, is the representation itself a visual representation, involving visual experiences? In *Consciousness Explained* (1991) Dennett favors the first hypothesis, which he sums up in his discussion of filling in: "The fundamental flaw in the idea of 'filling in' is that it suggests that the brain is providing something when in fact the brain is ignoring something" (p. 356).

We understand Dennett to mean that in the monocular condition the person may represent that there is a non-gappy object, say a vertical bar, in his visual field, but not because his brain generates a non-gappy *visual* representation of the vertical bar. In explicating his positive view on filling in, Dennett invites us to understand filling in of the blind spot by analogy to one's impression on walking into a room wallpapered with pictures of Marilyn Monroe:

> Consider how the brain must deal with wallpaper, for instance... Your brain just somehow represents *that* there are hundreds of identical Marilyns, and no matter how vivid your impression is that you see all that detail, the detail is in the world, not in your head. And no figment gets used up in rendering the seeming, for the seeming isn't rendered at all, not even as a bit map. (pp. 354–5)

If, as instructed, we are to apply this to the case of filling in of the blind spot, presumably Dennett's point is that no matter how vivid one's impression that one sees a solid bar, one's brain actually just represents that there is a solid bar. Dennett's claim, as he clarifies later, is that the brain ignores the absence of data from the region of the blind spot. In what follows, we shall show that, contrary to Dennett, the data strongly imply that at least some instances of filling in do indeed involve the brain "providing" something.

One preliminary semantic point should be made to forestall needless metaphysical tut-tutting. Hereafter, in discussing whether someone's perception of an object, say an apple, is filled in, we shall, *as a convenient shorthand*, talk about whether or not "the apple is filled in." In availing ourselves of this expedient, we do *not* suppose that there might be a little (literal) apple or (literal) picture of an apple in someone's head which is the thing that is filled in. Rather, we refer merely to some property of the brain's *visual* representation such that the perceiver sees a non-gappy apple.

Very crudely speaking, current neurobiological data suggest that when one sees an apple, the brain is in some state that can be described as representing an apple. This representation probably consists of a pattern of activity across some set of neurons, particularly those in the visual cortex, that have some specific configuration of synaptic weights and a specific profile of connectivity (P. M. Churchland 1989b). Given this general characterization of a representation, the question we want to address can now be rephrased: Does filling in an apple-representation consist in the visual cortex generating a representation which more closely resembles the standard case of an apple-representation of an apple in a peripheral visual field? Or does it consist, as Dennett (1991) suggests, in a nonvisual representation rather like one's nonvisual representation of the dog under the bed?

Our approach to these questions assumes that a priori reflection will have value mainly as a spur to empirical investigation, but not as a method that can be counted upon by itself to reveal any facts. Thought-experiments are no substitute for real experiments. To understand what is going on such that the blind spot is seen as filled in (non-gappy), it will be important to know more about the psychological and neurobiological parameters. In addition to exploring filling in of the blind spot, other versions of visual filling in, such as the filling in experienced by subjects

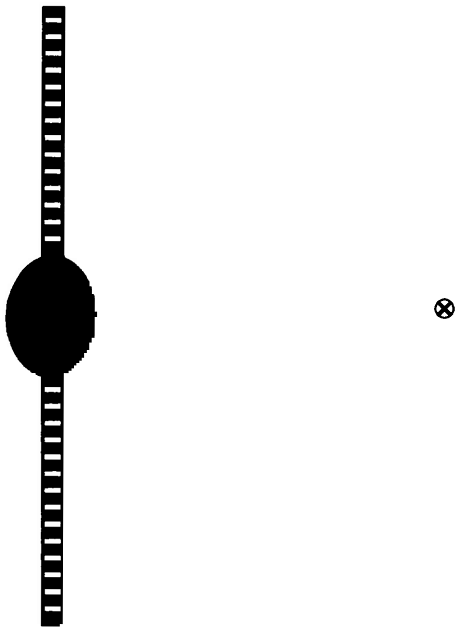

Figure 12.3
Subjects are presented with a display consisting of two vertical bar segments, separated by a gap of about 5 degrees, and this gap is positioned to coincide with the subject's blind spot. Fixation is to the right for left-eye viewing. Subjects report seeing an uninterrupted bar.

with cortical lesions, can also be studied. Although a more complete study would make an even wider sweep, embracing modalities other than vision, for reasons of space we narrow the discussion to visual filling in.

II. Psychophysical Data: The Blind Spot

To investigate the conditions of filling in, Ramachandran (1992) presented a variety of stimuli to subjects who were instructed to occlude one eye and fixate on a specified marker. Stimuli were then presented in various parts of the field in the region of the subject's blind spot. If a bar extends to the boundary on either side of the blind spot, but not across it, will the subject see it as complete or as having a gap (figure 12.3)? Subjects see it as complete. If, however, only the lower bar segment or the upper bar segment is presented alone, the subject does not see the bar

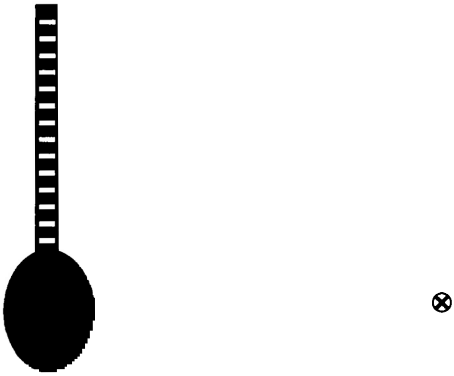

Figure 12.4
If only the upper segment of the bar is presented, subjects do not complete across the blind spot.

as filled in across the blind spot (figure 12.4). What happens when the upper bar and the lower bar are different colors; for example, upper, red; lower, green? Subjects still see the bar as complete, with extensions of both the red and green bar, but they do not see a border where the red and green meet, and hence they cannot say just where one color begins and the other leaves off. (For the explanation of nonperception of a border in terms of semisegregated pathways for functionally specific tasks, see Ramachandran 1992.)

Ramachandran also found that spokes extending to but not into the blind-spot boundary were filled in, demonstrating that filling in can be very complex. Suppose there is a kind of competition between completion of a black bar across the blind spot, and completion of an illusory contour lengthwise across the blind spot. Will the illusory contour or the real contour complete? Ramachandran discovered that in this test, the *illusory* contour typically completes (figure 12.5).

Ramachandran next explored the relation between subjective completion of a figure, and that figure's role in illusory motion (figure 12.6). The basic question is this: Does the brain treat a filled-in bar like a solid bar or like a gappy bar? In the control case, the upper gappy bar is replaced with the lower gappy bar (delay about 100 to 200 ms). Because the gap in the upper bar is offset with respect to the gap in the lower bar, subjects see illusory motion in a diagonal direction from left to right.

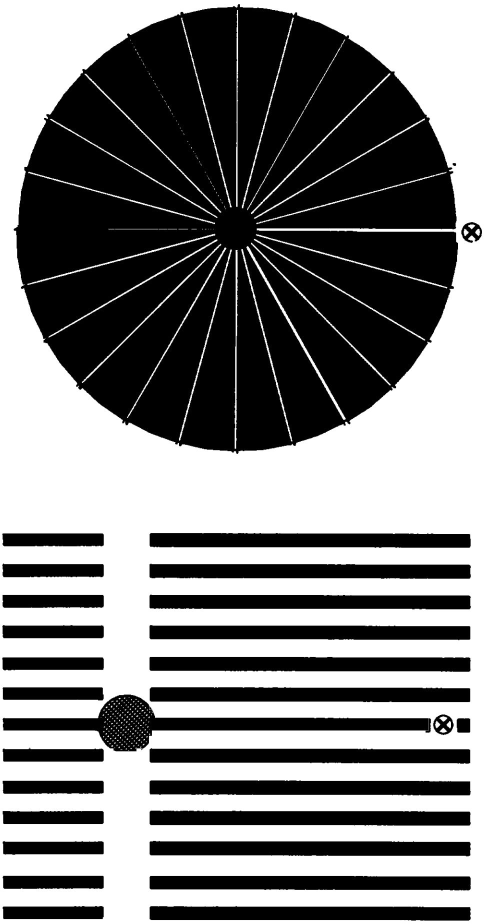

Figure 12.5
(a) Subjects reported perceptual completion of the spokes. (b) An illusory vertical strip was displayed so that a segment of the illusory contour fell on the blind spot. Subjects reported completion of the illusory strip rather than completion of the horizontal lines.

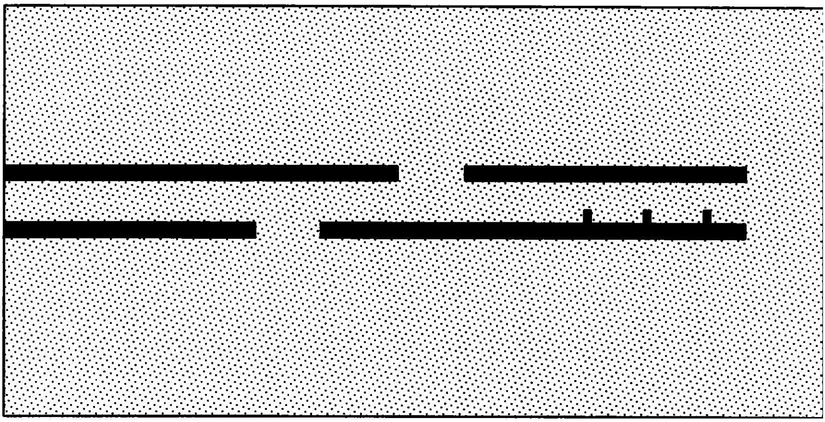

Figure 12.6
To generate illusory motion, the upper bar is replaced by the lower bar. When the gaps in both bars are located outside the blind spot, subjects see diagonal movement. When the gap in the upper bar coincides with the subject's blind spot, the movement appears to be vertical.

In the experimental (monocular) condition, the gap in the upper bar is positioned so that it falls in the subject's blind spot, and the subject sees a completed bar. Now when the upper bar is replaced by the lower bar to generate illusory motion, subjects see the bar moving vertically, *non-diagonally*, just as one does if a genuinely solid bar is replaced by the lower bar. This experiment shows that the brain treats a completed bar just at it treats a genuinely non-gappy bar in the perception of illusory motion.

According to Dennett's characterization of filling in (1991, p. 356), the brain follows the general principle that says, in effect, "just more of the same inside the blind spot as outside." Several of Ramachandran's results are relevant to this claim. If filling in is just a matter of continuing the pattern outside the blind spot, then in figure 12.7 subjects should see an uninterrupted string of red ovals, as a red oval fills the blank space where the blind spot is. In fact, however, subjects see an interrupted sequence; that is, they see two upper red ovals, two lower red ovals, and a white gap in between. In a different experiment, subjects are presented with a display of "bagels," with one bagel positioned so that its hole falls within the subject's blind spot (figure 12.8). The "more of the same" principle

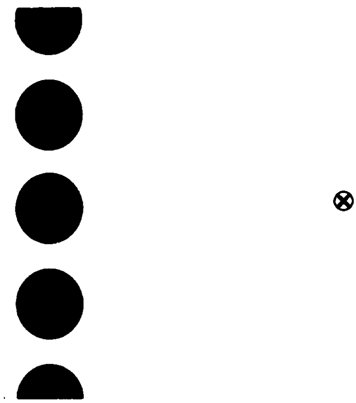

Figure 12.7
If the display is positioned so that a circle falls in the blind spot, subjects report a gap, not completion.

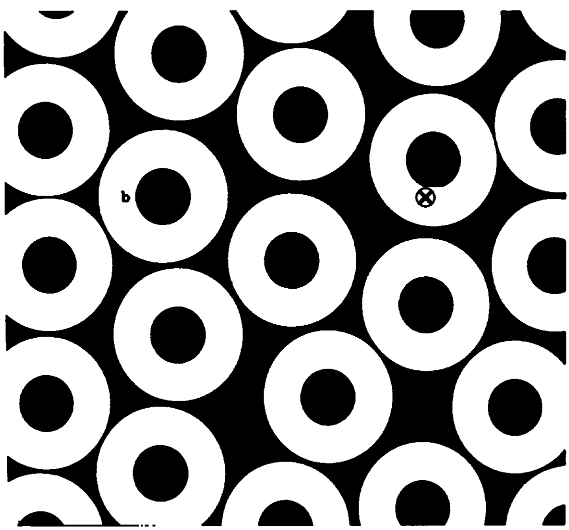

Figure 12.8
The display consists of yellow bagels and a fixation marker. The hole in one bagel (labeled *b*) coincides with the subject's blind spot. Subjects report seeing a yellow disk at this location, indicating that the yellow bagel is filled in.

presumably predicts that one will see only bagels in the display, as one apparently sees "more Marilyns." So the blind spot should not fill in with the color of its surrounding bagel. In fact, however, this is not what happens. Subjects see bagels everywhere, *save in the region of the blind spot*, where they see a disk, uniformly colored.

III. Psychophysical Data: Cortical Scotomata

A lesion to early areas of visual cortex (V1, V2; i.e., areas 17, 18) typically results in a blind area in the visual field of both eyes. The standard optometric test for determining scotomata consists in a flashing point of light in various locations of the visual field. Subjects are instructed to indicate, verbally or by pressing a button, when they see a flash. Using this method, the size and location of a field defect can be determined. Ramachandran (1992, 1993) explored the spatial and temporal characteristics of filling in of the scotoma in two patients (B. M. and J. R.).

B. M. had a right occipital-pole lesion caused by a penetrating skull fracture. He had a paracentral left hemifield scotoma, 6 × 6 degrees with clear margins. J.R. had a right occipital lesion caused by hemorrhage and a left visual field scotoma 12 degrees in width and 6 degrees in height. The locations of the lesions were determined by magnetic resonance (MR) scanning. Both patients were intelligent and otherwise normal neurologically. Vision was 20/20. B. M. was tested six months and J. R. eight months after the lesion events. Neither patient experienced his scotoma as a gap or hole in his visual field, but each was aware of the field defect. For example, each noticed some instances of "false" filling in of a real gap in an object. Additionally, they noticed that small, separable components of objects were sometimes unperceived, and noticed as missing. For example, one subject mistook the women's room for the men's room because the "Wo" of "Women's" fell into the scotoma. Seeing "men's," the subject walked directly in and quickly discovered his mistake.

In brief, the major findings of Ramachandran are as follows:

1. A 3-degree gap in a vertical line is completed across the scotoma, the completion taking about 6 seconds. The duration was determined by asking the patients to press a button when the line segment was completely filled in. Even with repeated trials, the latency remained the same.

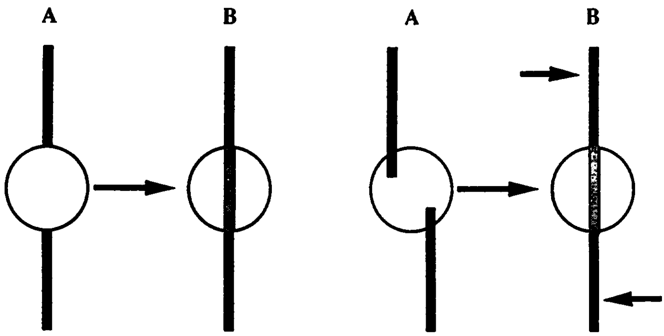

Figure 12.9
Schematic illustration of the stimuli shown to patients. The circle represents (roughly) the region of the patient's scotoma; fixation was approximately center field. (*Left*) Two bar segments were displayed on either side of the scotoma. The bar was vividly completed; the process of completion took about 6 seconds. (*Right*) The vertical bar segments were misaligned in the horizontal plane. After a few seconds of viewing, patients reported the lines moving toward each other until they became colinear. Then they gradually began to complete across the scotoma.

2. One patient (J. R.) reported that his perception of the filled-in line segment *persisted* for an average of 5.3 seconds after the upper and lower lines were turned off. The delay in completion as well as the persistence of "fill" is intriguing, and it is not seen in nontraumatic blind-spot filling-in.
3. When the top and bottom segments of the line were misaligned hori zontally by 2 degrees, both patients first reported seeing two misaligned segments separated by a gap. After observing this for a few seconds, they spontaneously reported that the upper and lower line segments began to drift toward each other, moving into alignment, then slowly (over a period of about 10 seconds) the line segments filled in to form a single line spanning the scotoma (figure 12.9). The realignment and visual completion took 6.8 seconds on average.
4. When viewing dynamic two-dimensional noise (e.g., "snow" on a television screen), one patient reported that the scotoma was first filled in with static (nonflickering) noise for 7 or 8 seconds before the random spots began to move and flicker. When the noise was composed of red pixels of randomly varying luminance, J. R. reported seeing the red color bleeding into the scotoma almost immediately, followed about 5 seconds later by the appearance of the dynamic texture.
5. When a vertical column of spots (periodicity greater than 2 degrees) was used instead of a solid line, both patients clearly saw a gap. When

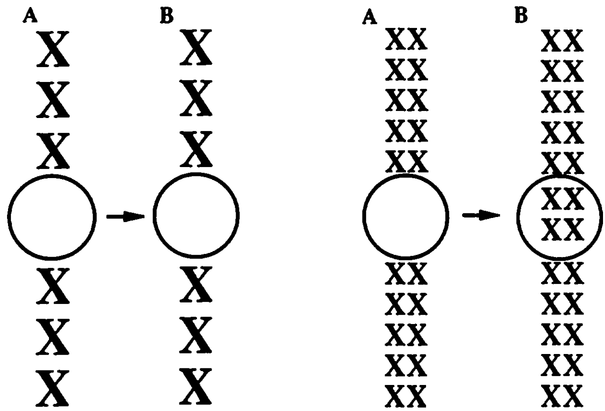

Figure 12.10
(*Left*) A column of large X's was not completed across the scotoma. (*Right*) A column of small X's did complete. If the column consisted of small horizontal line segments, the results were similar.

the spacing was reduced (periodicity less than 0.3 degree), patients reported seeing completion, across the scotoma, of a dotted line. These conditions were repeated using X's instead of spots, and the results were comparable (figure 12.10). Presented with a wavy, vertically oriented sinusoidal line (0.5 cycle per degree) with a gap matching the height of the patient's scotoma, both patients reported clearly seeing a non-gappy sinusoidally wavy line.

6. Each patient reported seeing illusory contours filled in across his scotoma. The experiment was similar to that performed with normal subjects (see again figure 12.5) save that the display was positioned so that the scotoma lined up with the gap in the stimuli. First, two horizontal line segments bordering the scotoma were presented, and as expected, they completed across the gap. Next, when an aligned array of horizontal lines was presented, the horizontal lines did not complete across the gap, and instead patients saw the vertical illusory strip complete across the scotoma.

7. Patients were presented with a checkerboard pattern , both fine (less than 0.3 degree)- and coarse (greater than 1.5 degree)-grained, which were readily filled in. When the checkerboard texture was subjected to counterphase flicker (7.5 Hz flicker; 0.6 check width), B. M. completed the flickering checks. J. R. however, reported that as soon as the pattern was made to flicker, he saw nonflickering stationary checks inside his scotoma, with the result that the margins of his scotoma became "en-

topically" visible. After about 8 seconds, J. R. saw the dynamic checks everywhere, including his scotoma.

8. To determine whether these filling-in effects might be seen in patients with laser-induced paracentral *retinal* scotomata, the tests were repeated on two such patients. Ramachandran found that (a) gaps in bars were not completed, (b) there was no motion or completion of misaligned bars, (c) coarse checkerboard patterns did not complete, (d) fine-grained 2-D random-dot textures were completed. This suggests that many of the completion effects are of cortical origin.

In the lesion studies, the time course for filling in, together with the subjects' reports, indicate that the completion is a visual phenomenon rather than a nonvisual judgment or representation. For example, when their spontaneous reports were tested with comments such as, "You mean you *think* that the checkerboard is uniform everywhere," the patients would respond with emphatic denials like, "Doctor, I don't merely *think* it is there; I *see* that it is there." Insofar as there is nothing in the visual stimulus corresponding to the filled-in perception, it is reasonable to infer, in contrast to Dennett, that the brain is "providing something," not merely "ignoring something." The visual character of the phenomenon also suggests that in looking for the neurobiological mechanism, visual cortex would be a reasonable place to start.

IV. Psychophysical Data: Artificial Scotomata

Ramachandran and Gregory (1991) discovered a species of filling in readily experienced by normal subjects, and conditions for which can easily be set up by anyone. The recipe is simple: adjust the television set to "snow" (twinkling pattern of dots); make a fixation point with a piece of tape, roughly in the middle of the screen; and place a square piece of gray paper, about 1 cm square and roughly isoluminant to the gray of the background, at a distance of about 8 cm from the fixation point (i.e., in peripheral vision). Both eyes may be open, and after about 10 seconds of viewing the fixation point, the square in peripheral vision vanishes completely. Thereafter, one sees a uniformly twinkling screen. Insofar as this paradigm yields filling in that is reminiscent of filling in of the blind spot and cortical scotoma, it can be described as inducing a kind of

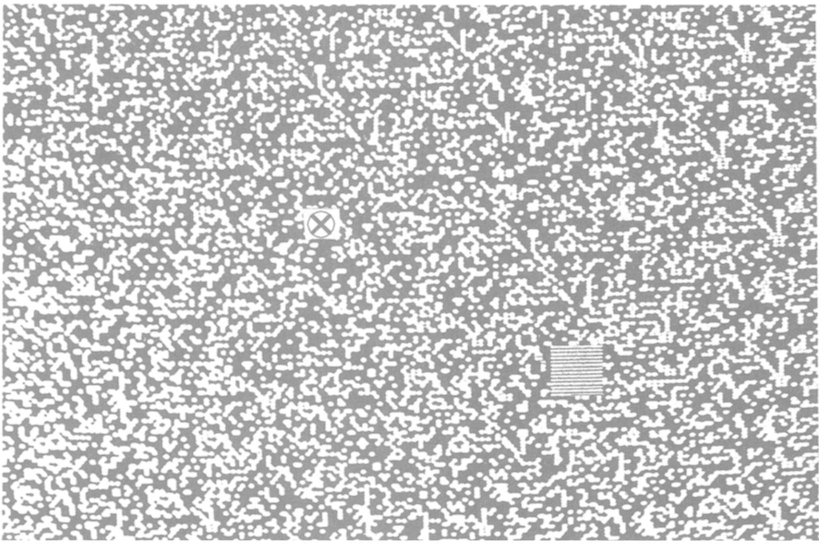

Figure 12.11
Display for artificial scotoma conditions consists of a background texture, a fixation point (the circle), and a small square segment in the peripheral field with a different texture, roughly isoluminant to the background.

artificial blind spot. Hence Ramachandran and Gregory called it an "artificial scotoma" (figure 12.11).

By using a computer to generate visual displays, many different arrangements of background texture and artificial scotomata can be investigated. In exploring the variety of conditions for filling in of an artificial scotoma, Ramachandran and Gregory found a number of striking results, several of which we briefly outline below:

1. Subjects tended to report filling in from outside the gray square to the inside, with a time scale of about 5 to 10 seconds.
2. Once subjects reported filling in to be complete, the background twinkles were then turned off. Subjects now reported that they *continued* to see twinkling in the "scotomic" square for about 3 to 4 seconds after the background twinkles had disappeared.
3. Suppose the background screen is pink and the twinkles are white. The "scotomic" square is, as before, gray, but within the square (which is now computer generated), dots are moving not randomly but coherently, left to right. Subjects report seeing a completely pink screen after about

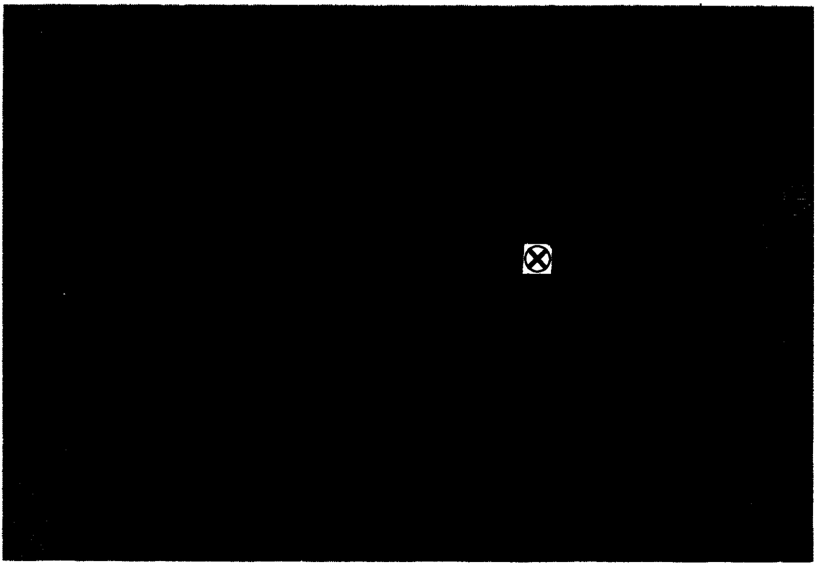

Figure 12.12
The filling in of the artificial scotoma can be complex. In this condition, subjects report that the text fills in.

5 to 10 seconds, but report that the dots in the square continue to move coherently from left to right. After a few more seconds, however, they report seeing uniformly random twinkles everywhere. Note that the artificial scotoma is in peripheral vision, where resolution is much poorer than in the foveal region (see again figure 12.2). The twinkles that are filled in look just the same as the twinkles elsewhere in the peripheral field.

4. If the screen is covered in text, the peripheral square comes to be filled in with text (figure 12.12).

5. If a smaller square with black borders is inscribed within the region of the gray square, subjects report that the inner area does not fill in with background texture.

Many other experiments in artificial scotomata are now underway in Ramachandran's laboratory, and those few cited here mark only the first pass in exploring a salient and intriguing visual phenomenon. For our purposes here, it is perhaps enough to note that, so far as we can determine, the results from artificial scotoma experiments do not confirm Dennett's hypothesis that, for phenomena such as filling in, "we can already be *quite* sure that the medium of representation is a version of something efficient, like color-by-numbers [which gives a single label to a whole region], not roughly continuous, like bit-mapping" (Dennett, 1991, p. 354).

V. Psychophysics and the Krauskopf Effect

Krauskopf (1963) discovered a remarkable filling-in phenomenon. In his setup, a green disk is superimposed on a larger orange disk. The inner boundary (between green and orange) is stabilized on the retina so that it remains on exactly the same retinal location no matter how the eyes jitter and saccade, but the outer boundary moves across the retina as the eyes jitter and saccade. After a few seconds of the image stabilization, the subject no longer sees a green disk; instead, the entire region is seen as uniformly orange—as filled in with the background color.

Using the Krauskopf image stabilization method to explore further aspects of the filling-in phenomenon, Thomas Piantanida and his colleagues have found more remarkable filling-in results. It is known that adaptation to yellow light alters a subject's sensitivity to a small flickering blue light; more exactly, flicker sensitivity is reduced in the presence of a yellow adapting background. Prima facie this is odd, given that "blue" cones are essentially insensitive to yellow light (it is the "red" and "green" cells that are sensitive to yellow light). Piantanida (1985) asked this question: Is blue flicker sensitivity the same if yellow adaptation is obtained by subjective *filling in of yellow* rather than by actual yellow light illuminating the retina?

To get a perception of yellow in an area where the retina was not actually illuminated with yellow light, Piantanida presented subjects with a yellow bagel, whose inner boundary was stabilized on the retina (using a dual Purkinje eye tracker) and whose outer boundary was not stabilized. The finding was that the yellow background achieved by image stabilization was *as effective* in reducing "blue" cone flicker sensitivity as an actual yellow stimulus. This probably means, therefore, that the reduction in flicker sensitivity as a function of perceived background is a cortical rather than a retinal effect. The most likely hypothesis is that cortical circumstances relevantly like those produced by retinal stimulation with yellow light are produced by yellow filling in, and hence the adaptation effects are comparable.

There is a further and quite stunning result reported by Crane and Piantanida (1983) that is especially relevant here. They presented subjects with a stimulus consisting of a green stripe adjacent to a red stripe, where the borders between them were stabilized, but the outside borders were not stabilized. After a few seconds, the colors began to fill in across the stabilized border. At this point, some observers described what they saw as a new and unnameable color that was somehow a mixture of red and green. Similar results were obtained with yellow and blue. Produced extraretinally, these visual perceptions of hitherto unperceived colors resulted from experimental manipulation of filling-in mechanisms—mechanisms that actively do something, as opposed to simply ignoring something.

Dennett says of the blind spot, "The area is simply neglected" (1991, p. 355). He says that "the brain doesn't have to 'fill in' for the blind spot, since the region in which the blind spot falls is already labeled (e.g., 'plaid' or 'Marilyns' or just 'more of the same')" (p. 355). Part of the trouble with Dennett's approach to the various filling-in phenomena is that he confidently prejudges what the neurobiological data at the cellular level will look like. Reasoning more like a computer engineer who knows a lot about the architectural details of the device in front of him than like a neurobiologist who realizes how much is still to be learned about the brain, Dennett jumps to conclusions about what the brain does not need to do, ought to do, and so forth.

In sections VI and VII below, we discuss neurobiological data that conflict with Dennett's claim that "There are no homunculi, as I have put it, who are supposed to 'care about' information arising from the part of the visual field covered by the blind spot, so when nothing arrives, there is no one to complain" (p. 357). And again: "The brain's motto for handling the blind spot could be: Ask me no questions and I'll tell you no lies" (p. 356). While Dennett's idea may seem to have some engineering plausibility, it is really a bit of a priori neurophysiology gone wrong. Biological solutions, alas, are not easily predicted from engineering considerations. What might, from our limited vantage point, have the earmarks of sound engineering strategy, is, as often as not, out of kilter with the way nature does it.

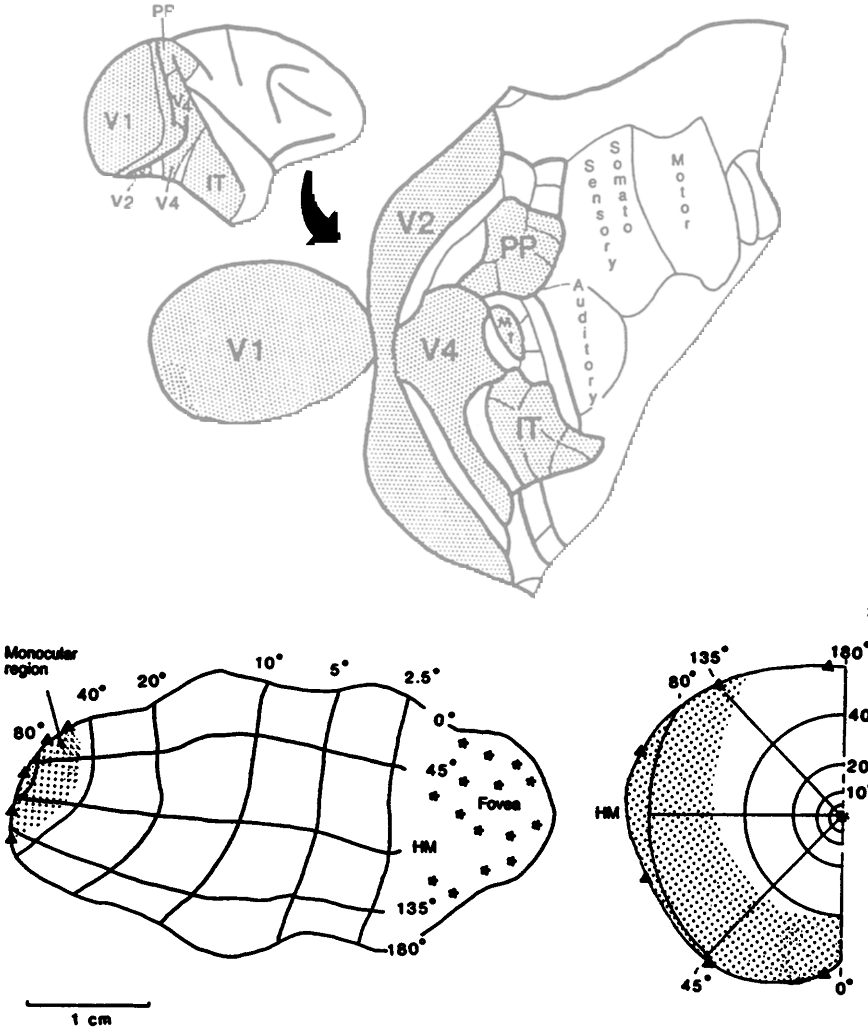

Figure 12.13
(*Above*) visual areas in the cerebral cortex of the macaque, as seen in a lateral view of the right hemisphere, and (arrow) in an unfolded two-dimensional map. The primary visual cortex (V1) is topographically organized. Lines of eccentricity (semicircles in the visual field drawing on lower right) map onto contours that run approximately vertically on the cortical map (*lower left*). Lines of constant polar angle (rays emanating from the center of gaze in the visual field) map onto contours that run approximately horizontally on the cortical map. The foveal representation (asterisks), corresponding to the central 2ϕ radius, occupies slightly more than ten percent of V1. The monocular region (stippled) in the visual field occupies a very small region of the cortical map. (Reproduced with permission from Van Essen and Anderson, 1990.)

VI. The Blind Spot and Cortical Physiology: The Gattass Effect

There are upward of twenty cortical visual areas in each hemisphere of monkeys, and probably at least that many in humans. Many of these areas are retinotopically mapped, in the sense that neighboring cells have neighboring receptive fields, that is, neighboring points in the visual field will be represented by neighboring cells in the cortex. In particular, visual area V1 has been extensively explored (figure 12.13). The receptive field size of V1 cells is about 2 to 3 degrees, and hence is much smaller than the size of the blind spot (about 6 × 4.5 degrees).

Ricardo Gattass and his colleagues (Fiorani et al., 1990; Gattass et al., 1992; Fiorani et al., 1992) were the first to try to answer the following question: How do V1 cells corresponding to the region of the blind spot for the right eye respond when the left eye is closed and a stimulus is presented to the open right eye (and vice versa)?

For ease of reference hereafter, by "Gattass condition" we denote the setup in which the experimenter records from single cells in V1 in the general area corresponding to the optic disk when the stimulus is presented to the contralateral (opposite side) eye. Call the V1 region corresponding to the optic disk of the contralateral eye, the "optic disk cortex" or ODC. The optic disk is that region of the retina where no transducers exist, corresponding to that part of the visual field where the blind spot resides. Remember that if a cortical region corresponds to the optic disk for the contralateral eye, it will correspond to *normal* retinal area for the ipsilateral (same side) eye. See figure 12.14 for projection patterns.

The seemingly obvious answer to Gattass' question—and the answer Gattass and his colleagues expected—is that the ODC cells will not respond in the monocular condition to stimuli presented in the contralateral blind spot. That is, one would predict that the cells in that region are responsive only to stimuli from the nonblind region of the ipsilateral eye. This is not what they found. Applying standard physiological mapping techniques to monkeys, and using the conventional bars of light as stimuli, they tested the responses of ODC cells (left hemisphere) with the left eye closed. As they moved the bar of light around and recorded from single cells, they found that neurons in the ODC area

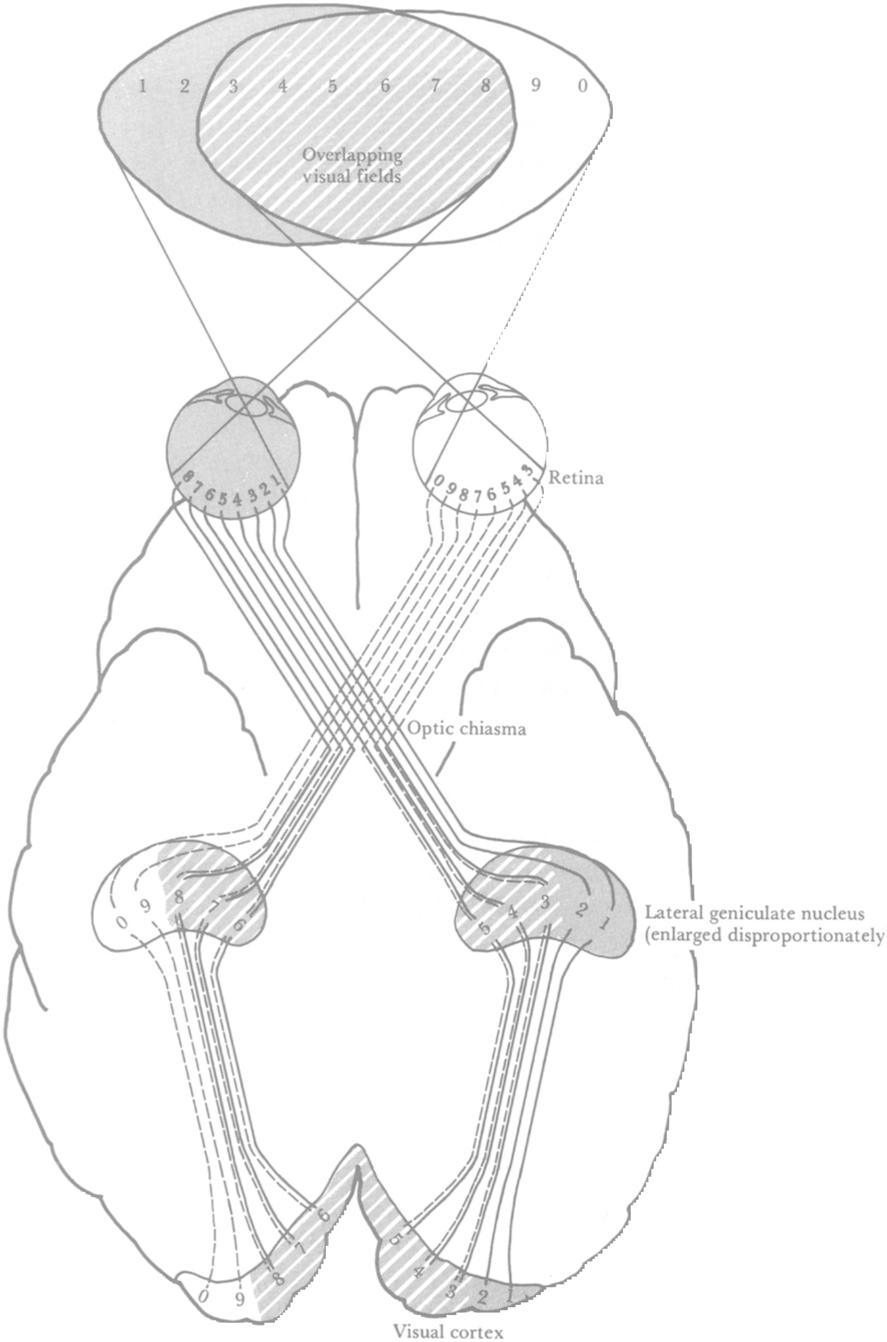
1 2 3 4 5 6 7 8 9 0
Overlapping visual fields
Retina
Optic chiasma
Lateral geniculate nucleus (enlarged disproportionately)
Visual cortex

responded very well. That is, cells corresponding to the blind spot gave consistent responses to a bar of light passing through the blind sector of the visual field. The response data did, however, show that the ODC was somewhat less neatly mapped by contralateral stimuli (i.e., in the blind spot) than by ipsilateral stimuli (i.e., in the nonblind field).

For some cells, an excitatory receptive field—presumably an interpolated receptive field—located *inside* the ODC could be specified. Exploring further, they found that sweeping bars on only one end of the blind spot yielded poor responses or none at all. In other cells, they discovered that the sum of two responses to two bar segments entering either end of the blind spot was comparable to the response for a single non-gappy bar. This indicates that some cells in the Gattass condition exhibit discontinuous receptive fields, presumably via interpolation signals from other neurons with neighboring receptive fields. To study the relevance of neighboring relations, Gattass and his colleagues masked the area immediately surrounding the optic disk during stimulus presentations to the blind spot region of the visual field. They discovered that responses of the ODC neurons were abolished in the masked condition (figures 12.15 and 12.16). Fifteen out of fourty-three neurons (mostly from layer 4ca) were found to exhibit interpolation properties across a region of the visual field at least three times the size of the classic receptive field.

VII. Artificial Scotomata and Cortical Physiology: The Gilbert Effect

How do cortical cells respond when their receptive field corresponds to the area of an *artificial* scotoma, such as the kind Ramachandran and Gregory studied? (For ease of reference, we shall hereafter call these

Figure 12.14
Schematic illustration of the projection pathways from the retina to the cortex, showing which parts of the visual field are represented in specific parts of the lateral geniculate nucleus (LGN) and the visual cortex. Note that the left hemifield projects to the right (contralateral), which in turn projects to the right hemisphere. The blind spot of the left eye corresponds approximately to the region coded as *3*, which is part of the central region where the fields of the two eyes overlap. By tracking *3* from the visual field, to the retina, to the LGN and the cortex, one can track the pathway for a particular stimulus in the blind region of the left eye. (Reproduced with permission from Lindsay and Norman, 1972.)

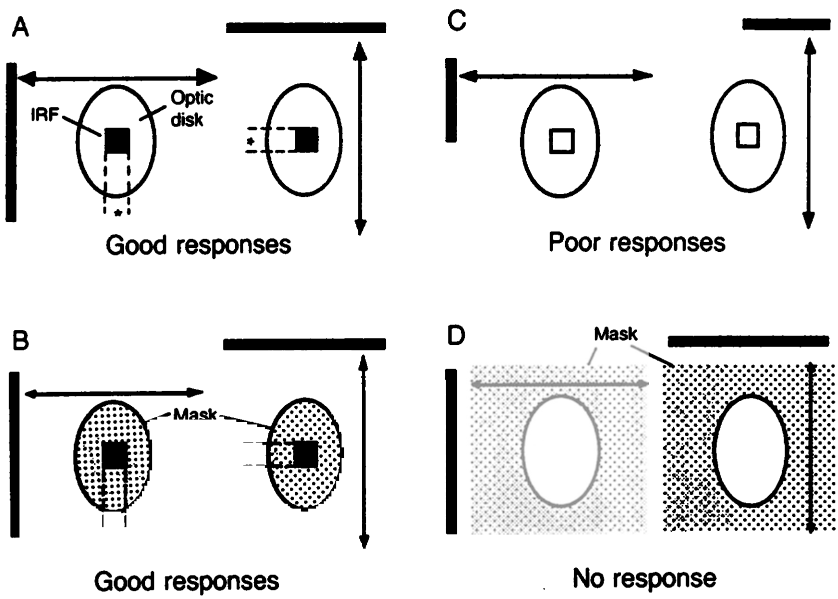

Figure 12.15
Summary of responses of cortical neurons in the optic disk cortical region to bars and masks of varying sizes. IRF: Interpolated receptive field; asterisks indicate locations where stimulation with long bars elicited responses. (After Fiorani et al. 1992.)

cortical neurons "artificial scotoma" or AS cells.) Or when the area from both retinas from which they receive projections is *lesioned*? (hereinafter, we shall call these cortical neurons "retinal lesion" or RL cells.) These questions have been addressed by Charles Gilbert and colleagues at Rockefeller University. Recording from V1 in monkeys, they discovered that the receptive fields of cortical cells surrounding the cortical RL cells expanded in several minutes so that collectively they covered that part of the visual field normally covered by the RL cortical cells (Gilbert and Wiesel, 1992).

A similar result was found in artificial scotoma experiments in cats (Pettet and Gilbert, 1991). The cortical cells in V1 surrounding the AS cortical cells very quickly expanded their receptive fields to include the area normally in the domain of the AS cells. The receptive field expansion was of the order of three- to fivefold. It was observed as soon as tests could be made (2 minutes), and it was reversible, in that once the ex-

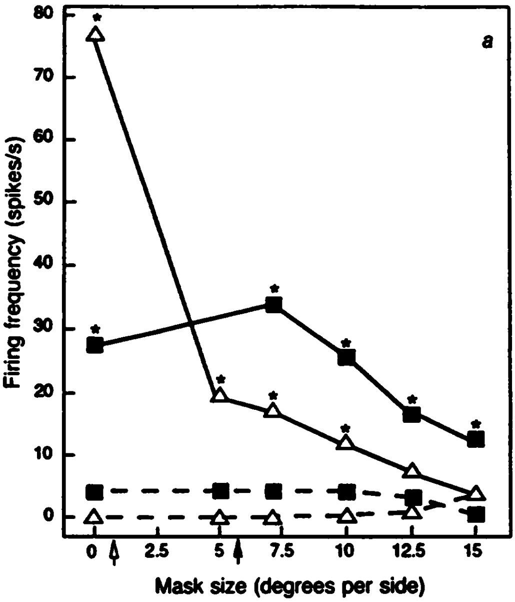

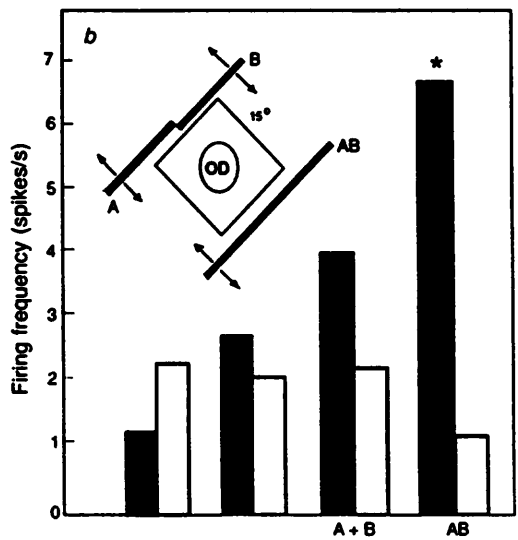

Figure 12.16
(a) Mean response rate of a V1 cell for ten presentations of the stimulus under different masking conditions. Triangles and filled squares linked by continuous lines show response to ipsilateral and contralateral eye, respectively, in paired trials. The lower (*dashed*) lines show the mean spontaneous activity where each eye is opened separately. The size of the ipsilateral classic receptive field is shown below by an outlined arrow, and the diameter of the optic disk (OD) by a filled arrow. (b) Black bars: mean response frequency of the same neuron to stimulation over a mask 15ϕ per side. Open bars show the mean spontaneous activity in paired trials without stimulation. (Reproduced with permission from Fiorani et al. 1992.)

perimental condition was removed and a standard, nonscotomatic stimulus was presented, normal mapping of cortical cells was restored. Although the neurobiological basis for this modification or interpolation in receptive field properties has not yet been determined, it is conjectured that lateral interactions within the cortex are probably crucial.

The Gattass effect, together with the Gilbert effect, is important evidence that the receptive fields of cortical cells are dynamic and can be modified on very short time scales. What precisely this means in terms of the neurobiological mechanisms of visual experience will require many more experiments. In any case, it is unlikely that the results are irrelevant to determining whether the brain merely ignores the blind spot or whether there is an active process related to filling in. As we try to track down the neurobiology of visual awareness, the discoveries in neuroscience are important clues to the nature of visual processing and to that component of the processing relevant to visual awareness.

Do the results from Gattass et al. and Gilbert et al. mean that, contrary to Dennett's assurances, filling in *is* rendered as a bit-map? No. The choices here are not exhausted by Dennett's alternatives "bit map or color-by-number." We suspect that neither the bit-map metaphor nor the color-by-number metaphor is even remotely adequate to the kind of representation and computation in nervous systems. Indeed, the very lability of a neuron's response properties and receptive field properties means that the bit-map metaphor is misleading. In order to understand more clearly how to interpret the results from Gattass and his colleagues, and from Gilbert and Wiesel, much more needs to be known about interpolation in neural networks and about the interaction of neurons within a mapped region and between adjacent regions. The fact is, very little is known at this point about the detailed nature of neural computation and representation, though we are at a stage where computer models highly constrained by neurobiological and psychophysical data can yield important clues (P. S. Churchland and Sejnowski, 1992).

VIII. Conclusion

In *Consciousness Explained*, Dennett brilliantly and quite properly debunks the idea that the brain contains a Cartesian Theater wherein images and the like are displayed. But the hypothesis that filling in (perceptual completion) may sometimes involve the brain's interpolating ("contributing something rather than ignoring something") certainly need have no truck whatever with Cartesian Theaters, either implicitly or explicitly, either metaphorically or literally, either sotto voce or viva voce. Given the data from psychophysics and neurophysiology, we hypothesize that (a) the brain has mechanisms for interpolation, some of which may operate early in visual processing; (b) brains sometimes visually represent completions, including quite complex completions; and (c) such representation probably involves those interpolation mechanisms.

How did Dennett come to embrace a conclusion so manifestly contrary to the data, some of which were readily available when his book was published? And why does "filling in" play such an important role in *Consciousness Explained*? According to our analysis, the background derives from the background behaviorist ideology that is endemic to Dennett's work from the very beginning—from his first book, *Content and Consciousness* (1969), through *Brainstorms* (1978), *Elbow Room* (1984), *The Intentional Stance* (1987), and *Consciousness Explained* (1991).

Simplified, the heart of Dennett's behaviorism is this: the conceptual framework of the mental does not denote anything real in the brain. The importance of the framework derives not from its description of neural or any other reality; rather, it is an organizing instrument that allows us to do fairly well in explaining and predicting one another's behavior, the literal unreality of qualia, and so forth, notwithstanding. How is it that the framework manages to be a useful instrument, despite the unreality of its categories? Because, according to Dennett, even though there is nothing *really* in the brain that corresponds to visual awareness of red, there is *something or other* in the brain which, luckily enough, allows us to get on pretty well in making sense of people's behavior on the pretense, as it were, that the brain really does have states corresponding to awareness of red. As for filling in, Dennett's rhetorical strategy hoists it as *paradigmatic* of a mental thing that we mistakenly assume to be real.

Dennett's discussions regarding the dubiousness of selected old-time intuitions often fall upon receptive ears because some categories such as "the will" and "the soul" probably do not in fact correspond to anything real, and because neuroscience is bound to teach us many surprising things about the mental, including that some of our fundamental categories can be improved upon. The sweeping behaviorism-instrumentalism, however, does not follow from these observations about the revisability of psychological concepts—nor even from the eliminability by cognitive neuroscience of *some* concepts that turn out to be the psychological counterpart of "phlogiston," "impetus," and "natural place." Thus one may readily concur that qualia cannot be little pictures displayed in the brain's Cartesian Theater, and that the self is not a little person tucked away in the folds of frontal cortex. These debunking treats are, however, just the teaspoon of sugar that makes the medicine go down. And the medicine, make no mistake, is behaviorism. The elixir is *Gilbert Ryle's ghost-be-gone* (Ryle, 1949). Taken regularly, it is supposed to prevent the outbreak of mental realism. Drawing on artificial intelligence's conceptual repertoire of the "virtual machine," Dennett has systematically argued *against* the neural reality, and *for* the mere instrumental utility, of mental categories generally. Dennett's engaging exposition and brilliantly inventive metaphors tend to mask the fact that this less palatable message is indeed the main message (see also McCauley, 1996).

This brief excursion through Dennett's behaviorism and instrumentalism may help explain why he is found defending assorted theses that are highly implausible from a scientific perspective: the brain does not fill in; there is nothing whatever (no "fact of the matter") to distinguish between a misperception and a misrecollection; there is no time before which one is not aware of, say, a sound, and after which one is aware; human consciousness is a virtual machine that comes into being as humans learn to talk to themselves; and so forth (Dennett 1991).

Scientific realism (P. M. Churchland, 1979, 1989), in contrast to Dennett's instrumentalism, proposes that we determine by empirical means—by converging research from experimental psychology, neuropsychology, and neuroscience—what hypotheses are probably true, and hence what categories truly apply to the mind-brain. Some current categories may be largely correct, for example, "visual perception"; some, for example, "memory," "attention," and "consciousness," appear to be subdividing, budding, and regrouping; and some may be replaced entirely by high-level categories that are more empirically adequate. At this stage, it is reasonable to consider sensory experiences to be real states of the brain, states whose neurobiological properties will be discovered as cognitive neuroscience proceeds (P. S. Churchland, 1986, 1988; P. M. Churchland, 1989a).

Perhaps Dennett's main achievement consists in showing the Cartesian dangers waiting to ensnare those who refer to perceptual filling in by means of the expression "filling in." If so, then the achievement is primarily semantic, not empirical. Furthermore, his aim could be satisfied merely by instructing us of the dangers, without requiring also that the very description "filling in" be expunged as untrue of what goes on in the brain. In any case, one might well wonder whether Dennett overestimates the naiveté among scientists. To judge from the literature, those who scientifically study perceptual completion phenomena understand perfectly well that filling in involves no Cartesian Theaters, ghosts, paint, little pictures, putty knives, or homunculi. At the very least, they are no more addled by metaphor than is Dennett when he refers to the brain as "editing multiple drafts." Taken as a linguistic prohibition rather than an empirical hypothesis about the mind-brain, Dennett's thesis that "the brain does not fill in" sound uncomfortably like a quirky edict of the 'word-police.'

Acknowledgments

For helpful discussion at many stages, we thank Paul Churchland, Francis Crick, Peter Dyan, Bo Dahlbom, Ricardo Gattass, Read Montague, Diane Rogers-Ramachandran, Adina Roskies, and Oron Shagrir.