Martinez-Conde, Macknik, Martinez, Alonso & Tse (Eds.)
Progress in Brain Research, Vol. 155
ISSN 0079-6123

CHAPTER 6

In honour of Lothar Spillmann — Filling-in, wiggly lines, adaptation, and aftereffects

Stuart Anstis*

Department of Psychology, UCSD, 9500 Gilman Drive, La Jolla, CA 92093-0109, USA

Abstract: I have studied a number of visual phenomena that Lothar Spillmann has already elucidated. These include:

Neon spreading: when a small red cross is superimposed on intersecting black lines, the red cross seems to spread out into an illusory disk. Unlike the Hermann grid, neon spreading is relatively unaffected when the black lines are curved or wiggly. This suggests that the Hermann grid, but not neon spreading, involves long-range interactions. Neon spreading can be shown in random-dot patterns, even without intersections. It is strongest when the red crosses are equiluminous with the gray background.

Adaptation, aftereffects, and filling-in: direct and induced aftereffects of color, motion, and dimming. Artificial scotomata and filling-in: the "dam" theory is false. Staring at wiggly lines or irregularly scattered dots makes them gradually appear straighter, or more regularly spaced. I present evidence that irregularity is actually a visual dimension to which the visual system can adapt.

Conjectures on the nature of peripheral fading and of motion-induced blindness.

Some failed experiments on correlated visual inputs and cortical plasticity.

Keywords: adaptation; aftereffects; afterimages; color induction; filling-in; illusions

Introduction

It is impossible to summarize Lothar Spillmann's contributions to visual psychophysics because he has studied just about everything. If he has not studied it, it is not psychophysics. I shall just discuss some random samples taken from his formidable body of works on vision. The topics I have picked include the Hermann grid, neon spreading, filling-in and aftereffects, and visual plasticity. Note that many of the illusions described here are beautifully illustrated on the web page of Lothar's colleague Michael Bach at http://www.michaelbach.de/ot/

Long- and short-range interactions: Hermann's grid vs. neon spreading

Hermann grid. Spillmann has always been interested in the relationships between long- and short-range interactions in vision (Spillmann and Werner, 1996; Spillmann, 1999). A case in point is the Hermann-grid illusion (Hermann, 1870; Spillmann, 1971, 1994; Spillman and Levine, 1971; Oehler and Spillmann, 1981), which has long been regarded as a short-range process but has now been shown to require long-range processes as well (Geier et al., 2004). In the Hermann grid, illusory dark spots or blobs can be seen at every street crossing, except for the ones that are being directly fixated. A stronger version, known as the scintillating grid (Schrauf et al., 1997; Ninio and Stevens, 2000; Schrauf and Spillmann, 2000),

---

*Corresponding author. Tel.: +1-858-534-5456;
E-mail: sanstis@ucsd.edu

DOI: 10.1016/S0079-6123(06)55006-X

has a small disk at each intersection. This produces a smaller but much darker and more vivid illusory point. Both the Hermann grid and the scintillating grid work equally well in reversed contrast, with black stripes on a white ground.

The standard, short-range explanation comes from Baumgartner (1960). He suggested that an on-center retinal ganglion cell could be positioned by chance at an intersection, in which case it would have four bright regions in its inhibitory surround, one from each street, and these would reduce its response. A ganglion cell looking at a street would have only two inhibitory regions, so it would respond more strongly. A fixated intersection falls on the fovea, where the receptive fields are so small that it would make no difference whether or not it fell on an intersection. In fact, Spillmann (1994) and Ransom-Hogg and Spillmann (1980) measured the stripe widths that gave the maximum illusion at different eccentricities in order to determine the size of human ‘‘perceptive fields.’’

This explanation fails to explain why global factors are important. Wolfe (1984) pointed out that Baumgartner's model is local in nature, since it relies on cells with concentric on-off or off-on receptive fields. This model predicts that the magnitude of the illusion at a given intersection should be the same whether that intersection is viewed in isolation or in conjunction with other intersections in a grid. However, Wolfe showed that illusion magnitude grows with the number of intersections and that this growth is seen when the intersections are arranged in an orderly grid but not when they are placed irregularly. These results rule out any purely local model for the Hermann-grid illusion. Global factors must be involved. Geier et al. (2004) decisively overthrew the Baumgartner model by imparting a slight sinusoidal curvature to the lines. When the lines are straight the illusion is visible, but as soon as the lines become curved the illusion vanishes. The same distortions applied to the scintillating grid made the scintillations disappear. This implies that the Hermann grid and the scintillating grid both depend upon long-range interactions, probably operating along the length of the lines (see Fig. 1).

Neon spreading. Spillmann has also studied the neon spreading that can be seen at the intersection of two thin black lines (Bressan et al., 1997). A red + sign superimposed on the intersection appears to spread out into a pink disk, provided that the black lines are continuous with, and aligned with, the red lines (Redies and Spillmann, 1981; Spillmann and Redies, 1981; Redies et al., 1984; Kitaoka et al., 2001). Don Macleod and I wondered whether neon spreading, like the Hermann grid, would vanish for curved lines. If so, neon spreading would also depend upon long-range global interactions, and not merely upon local factors. Accordingly we (he) wrote a program that could apply any desired curvature to a neon-spreading lattice of black lines. *Result*: Curving the lines did *not* reduce the neon spreading, in sharp contrast to Geier's results with the Hermann grid. This suggests that neon spreading is a *local*, short-range affair.

Fig. 2 shows that neon spreading is strongest when the red crosses are equiluminous with the surround. In Fig. 2, the gray background is swept from dark on the left to light on the right, while the red crosses are swept from darkest at the bottom to lightest at the top. A glance at Fig. 2a shows that neon spreading is strongest along a positive diagonal where the luminances of the colored crosses and the gray background are equal.

In that case, what is the minimum stimulus that neon spreading requires? My own observations suggest, not much. It is well known that a square lattice of thin black lines on a white surround gives strong neon spreading when the intersections are replaced with red. But I also produced neon spreading in sparse, stationary random black dots scattered on a white surround, simply by coloring a ring-shaped subset of the black dots red (not illustrated). The ring was then moved around, but the red/black dots defining it remained stationary, merely turning red when they lay within the annular region that defined the moving ring and returning to black when they did not. Result: observers reported a pink neon annulus moving around across a stationary random-dot field. The neon effect was much stronger when the ring moved than when it was stationary. This shows that neon spreading is not necessarily dependent upon geometrical features such as intersections. It merely needs to replace black regions that lie on a white ground.

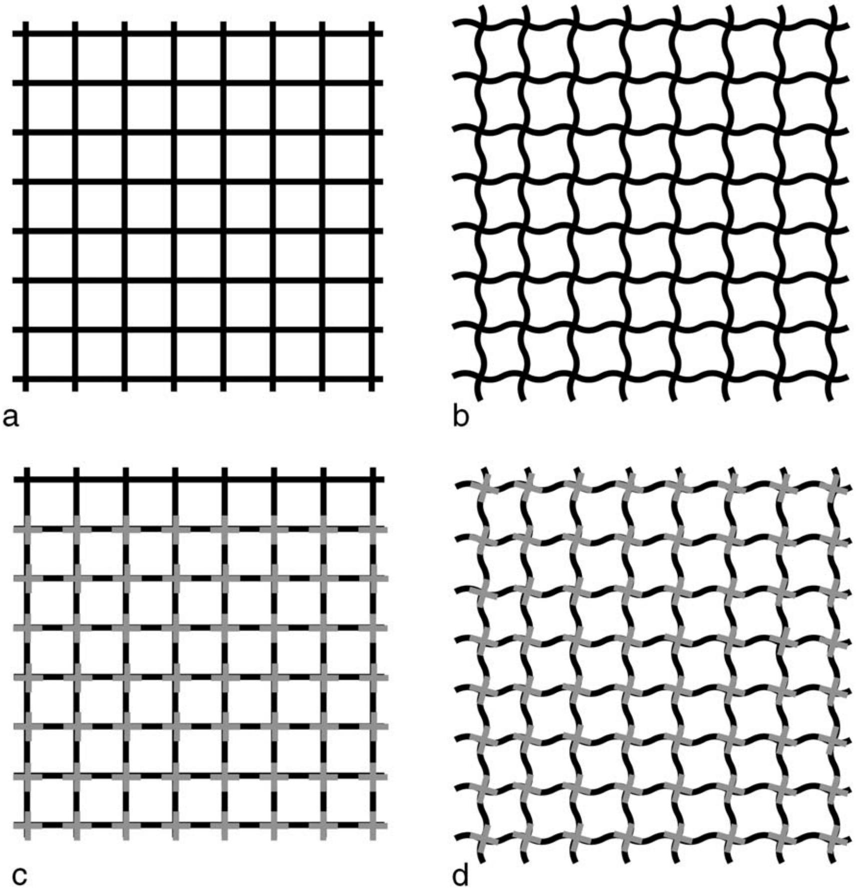

Fig. 1. (a) Hermann grid …(b) is abolished by curving lines (Geier et al., 2004). (c) Neon spreading…(d) is unaffected by curving lines.

These results point to some low-spatial-frequency visual pathway that blurs the seen colors and spreads them outside the confines of the thin lines. Could neon spreading simply be a by-product of the famously low acuity of the chromatic pathways? (Kelly, 1983). This cannot be the whole story since neon spreading also works for gray (Bressan et al., 1997). It might be that the beautiful watercolor effect discovered by Pinna et al. (2001, 2003) is an extreme case of neon spreading.

Peripheral fading

Lothar Spillmann has always been fascinated by the fact that strict fixation can make peripherally viewed stimuli fade out and disappear from view. Here are some examples of peripheral fading.

1. Troxler fading of a luminance-defined object, such as a black or white disk on a gray surround (Troxler, 1804).
2. A window of drifting dense random dots embedded in a field of twinkling dynamic noise gradually fades from view and disappears (Anstis, 1989). When all motion is subsequently stopped, the window shows a negative aftereffect of motion (Figs. 3a, b).
3. A small, peripheral gray patch embedded in a field of twinkling dynamic noise also gradually fades from view (Ramachandran and Gregory, 1991; Spillmann and Kurtenbach,

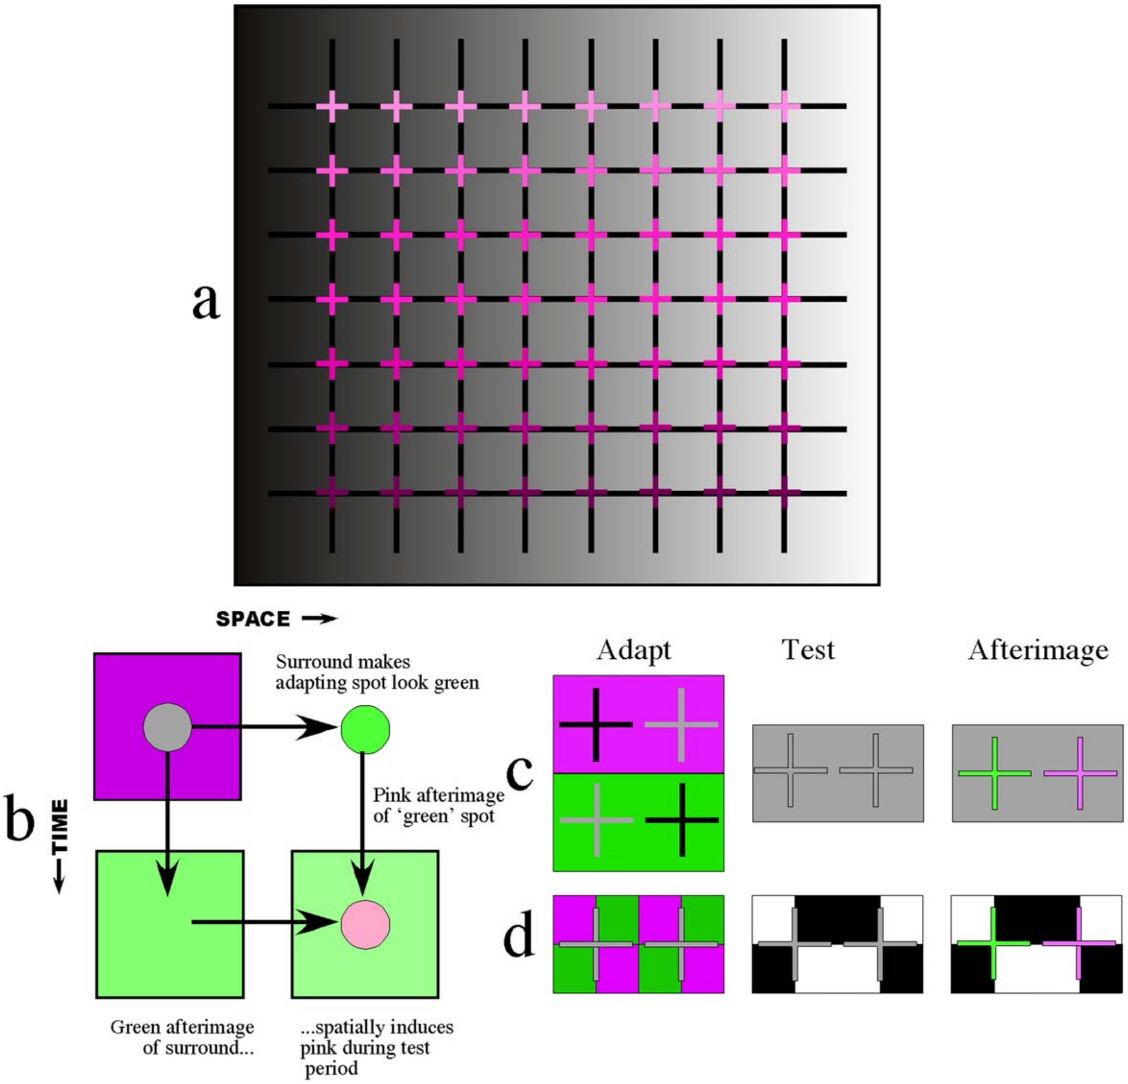

Fig. 2. (a) Neon spreading is strongest when the colored crosses are equiluminous with the gray surround. (b) Two possible neural routes for induced afterimages: (i, top right) surround spatially induces subjective green into during adaptation, then greenish spot has pink afterimage; (ii, bottom left) magenta surround has green afterimage, which spatially induces pink into the post during the test period. (c) shows existence of (i): following adaptation to alternating black and gray cross on magenta and green surrounds, a gray test field shows afterimages in which the left cross looks greenish and the right cross looks pinkish. (d) shows existence of (ii): following adaptation to gray crosses lying on green and magenta quadrants, the white test quadrants show afterimages which induce the left cross to look greenish and the right cross to look pinkish. (After Anstis et al., 1978.)

1992). When the field switches to a uniform gray test field, a twinkling aftereffect, resembling the original twinkling dots, fills the gray patch (Figs. 3c, d).

4. A gray patch, or a patch of one kind of static texture, embedded in another kind of static texture, will also fade out and disappear (Spillmann, 2003). The more salient the texture, the longer it takes to disappear (Sturzel and Spillmann, 2001).
5. A flickering peripherally viewed spot remains visible, but the perceived amplitude of its

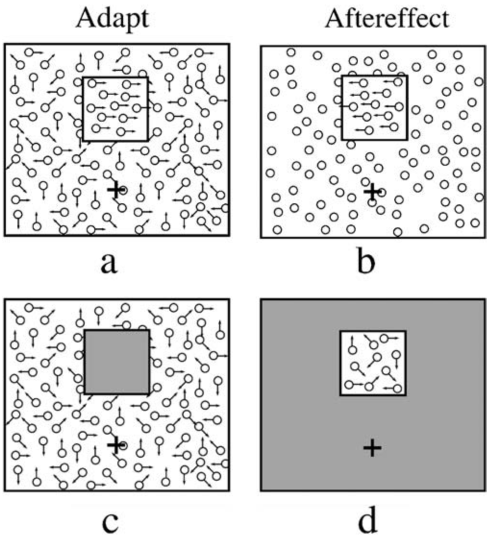

Fig. 3. (a) A window filled with drifting dots, embedded in dynamic noise, rapidly disappears from view and (b) gives a negative motion aftereffect on a stationary test field (Anstis, 1989). (c) A gray patch also disappears and (d) gives a twinkling aftereffect (Ramachandran and Gregory, 1991).

flicker falls steadily over time (Schieting and Spillmann, 1987; Anstis, 1996).

6. Peripherally viewed motion can appear to slow down and stop (Campbell and Maffei, 1979; Hunzelmann and Spillman, 1984).
7. A very blurred colored spot on an equiluminous gray surround gradually fades from view. On a gray test field an afterimage in the complementary color is visible. Jeremy Hinton has a nice demonstration on Michael Bach's web page. Twelve identical small blurred purple spots are arranged around a clock face, and the observer fixates in the middle. The spot at 1 o'clock is briefly turned off and then turned back on. Then the spot at 2 o'clock similarly disappears briefly, then the spot at 3 o'clock, and so on. At first one sees a "traveling gap," sometimes knows as omega movement (Tyler, 1973). But after a few seconds all the purple spots fade out from view and one simply perceives a bright green spot — an afterimage — running around clockwise.
8. A spectacular new illusion of peripheral fading is known as "motion-induced blindness" (Bonneh et al., 2001).

Filling-in and aftereffects: Spillmann's theory

Peripherally viewed stimuli often fade out from view during strict fixation, as listed above. For instance, if a large surround is filled with twinkling dynamic noise, a simple small gray square (Ramachandran and Gregory, 1991) or a static window filled with drifting random dots (Anstis, 1989) will disappear from view within ~10 s. Spillmann and de Weerd (2003) have suggested a *two-stage theory* to explain all such examples of delayed peripheral fading. First, there is a slow adaptation of the figure's boundary representation that normally keeps the figure perceptually segregated from its surround, and then a fast filling-in or interpolation process takes place in which the area previously occupied by the figure becomes invaded by the surround. De Weerd et al. (1998) measured the time for filling-in of a gray square on a black background filled with small white vertical bars that moved dynamically, a form of anisotropic dynamic noise. They used various sizes and eccentricities of squares, and based upon estimates of cortical magnification, they found that the time for filling-in was linearly related to the total contour length of the square's projection upon the visual cortex, rather than its retinal image size. This is consistent with Spillmann and de Weerd's theory. I shall characterize (or caricature) Spillmann and de Weerd's theory as the "dam" theory. The boundary of a peripheral gray patch walls or dams it off from an ocean of twinkling dots. Adaptation slowly erodes the dam, and twinkling water then rushes in quickly. This fits the experimental facts. However, Ramachandran and Gregory applied a dam theory to the subsequent aftereffect, such that when the water is drained away, some twinkling water remains briefly trapped inside the dam. This story predicts that the water inside the dam — the aftereffect — should be the same color or texture as the adapting ocean. But I find that it is not. The twinkle aftereffect has a fixed grain size (spatial frequency), regardless of the adapting grain size!

Ramachandran and Gregory (1991) attribute their twinkling aftereffect to a process of interpolation that actively fills-in the gray test square with the twinkling dots from the adapting surround. However, Tyler and Hardage (1998) disagree, and my own observations also suggest otherwise. Like Ramachandran et al., (1991) I used a gray patch set in a field of twinkling dots, but I also systematically varied the grain size of the adapting random dots over an eightfold range, with dots ranging from 0.1° to 0.8° in diameter. When observers were then exposed to a uniform gray test field, they were able to match the apparent grain of the twinkling aftereffect by means of an adjustable random-dot field. If some active filling-in were responsible, the aftereffect grain size should match the adapting grain size. In fact, however, I found that the aftereffect was always matched to the *same* grain size of 0.1°, whatever the adapting grain be. Since the aftereffect remained constant even when the adapting field changed, it was clearly not the result of a simple filling-in process. But in that case, what was it? (see Fig. 4).

A clue comes from "induced afterimages of color" (Anstis et al., 1978). If you adapt to an equiluminous gray cross in a very large green surround and then switch to a uniform gray test field, the expected afterimage might be a gray cross in a pink surround. But in fact, the afterimage is of a strongly green cross in a neutral gray surround! Thus, the cross afterimage matched the adapting surround. But this is not a simple fill-in. Instead, there are two possibilities (Fig. 2b). Either the adapting surround spatially induced an apparent pink into the cross during the adapting period, and this apparent pink was followed by its own afterimage in the cross, or alternatively, the adapting surround was followed by its own pink afterimage, which then spatially induced green into the cross during the test period. In the first case, simultaneous contrast (spatial induction) precedes successive contrast (afterimage). In the second case, successive contrast precedes simultaneous contrast. Many years ago (Anstis et al., 1978) we showed that both processes occur and can be elicited independently. We arranged for the background to switch every few seconds between two complementary colors, green and magenta. Two thin crosses lay side by side (Fig. 2c). When the surround was green, the left-hand cross was an equiluminous gray. By Grassman's (1853) third law, this induced a strong apparent pink into the cross. The right-hand cross was black, and only a minimum amount of pink was induced into it. Conversely, when the surround was magenta, the left-hand cross was black and the right-hand cross was an equiluminous gray, and looked apparently greenish. Following an adaptation period of 30–60 s, a uniform gray test field was presented. Two cross-shaped afterimages were visible. The left-hand afterimage looked green, and the right-hand afterimage looked pink. This shows that the subjective colors induced into the crosses during the adaptation period could generate their own afterimages. Note that the alternation of the complementary colors green and magenta in the adapting surround would cancel out and produce no net colored afterimage in the test surround.

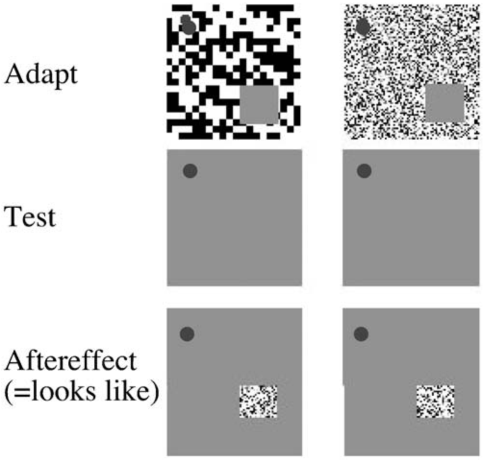

Fig. 4. Breakdown of the dam theory. If one fixates the black dot, the peripheral gray square soon disappears (top row), whatever the grain size of the noise. This fits the dam theory. When a gray test field is substituted (middle row) the previous location of the gray square is filled with a twinkling aftereffect (bottom row). However, this is not the simple fill-in that a dam theory might predict because the grain size of the aftereffect is constant and unrelated to the adapting grain size.

In our second experiment, the adapting fields surrounding each cross were divided into green and magenta quadrants of equal sizes (Fig. 2d). The crosses always looked neutral gray because each was bordered by equal amounts of green and magenta, during the adapting period. But during the test period the surround quadrants were made black and white. The black quadrants showed little or no afterimage, but on the white quadrants strong negative afterimages were seen, which spatially induced secondary pink and green afterimages into the crosses. As before, the left-hand afterimage looked green and the right-hand afterimage looked pink, but now the reasons were different. This result shows that the surround afterimages could spatially induce subjective colors into the crosses during the test period, but not the adaptation period. Thus in our first experiment, simultaneous contrast preceded successive contrast, while in our second experiment, successive contrast preceded simultaneous contrast. An alternative formulation is that the visual system can adapt to color ratios that can be expressed as edge-redder-on-left and edge-greener-on-right. This description includes red/white, red/green, and white/green edges. Incidentally, in these experiments an outline of the crosses was included in the test field because afterimages are easier to see when they are outlined (Daw, 1962). I have also found similar interactions between simultaneous induction and successive aftereffects, both for motion (Anstis and Reinhardt-Rutland, 1976) and for adaptation to gradual change of luminance (Anstis, 1979), which I shall not describe here. For a useful review of color induction, see Zaidi (1999). For filling-in, see Pessoa and de Weerd, (2003).

I conclude that in these experiments on colored afterimages, and in the disappearance of small patches superimposed on texture (Anstis, 1989; Ramachandran and Gregory, 1991; Spillmann and Kurtenbach, 1992; Sturzel and Spillmann, 2001; Spillmann, 2003), there is a complex interplay between processes of simultaneous and successive contrast. The aftereffects were *not* a simple fill-in from the adapting surrounds.

Motion-induced blindness

Bonneh et al. (2001) have discovered a dramatic example of peripheral fading. Three small stationary yellow spots forming a triangle are displayed on a monitor screen, against a background of small dark blue spots that rotate around the center of the screen or else fly around randomly like a swarm of midges. If one gazes at the center of the triangle, the yellow spots dramatically disappear and reappear. Although the conditions are not too critical, the effect is strongest for small, high-contrast stationary yellow dots against numerous, high-contrast blue dots in rapid motion. The blue surround spots can be flickering instead of moving.

Fig. 5 shows a simplified version of their stimulus, comprising an array of stationary dark blue flickering spots. Each spot flickers independently between 100 different random luminance levels, retaining always the same blue hue. Three of the spots are static yellow instead of blue. When the observer looks at the middle of this triangle of yellow spots, the yellow spots seem to disappear

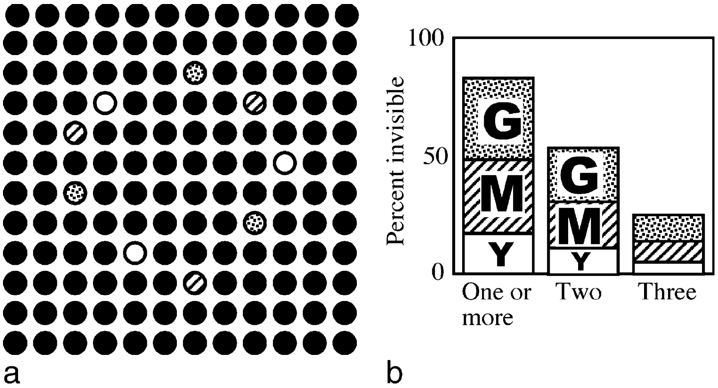

Fig. 5. (a) All spots flicker except the three yellow (unfilled) spots, which seem to disappear dramatically. If observers attend to the three flickering magenta (striped) or flickering green (spotted) dots, these also disappear. (b) Percentage of accumulated invisibility period for the disappearance of one or more spots, exactly two spots, and exactly three spots. The spots were invisible for about 50% of the time. The flickering green and magenta spots disappeared even more than the static yellow ones.

and reappear. Bonneh et al. refer to this as "motion-induced blindness." However, in my hypothesis motion does not actually *induce* anything. Instead, I conjecture that:

1. *All peripherally viewed targets tend to disappear over time.* This conjecture applies to all the peripheral stimuli listed above, not just to motion-induced blindness.
2. *Changing (*i.e., *flickering or moving) targets are more resistant than static targets to disappearing.* This does not mean that flickering or moving objects are a proof against disappearance: they are not (Campbell and Maffei, 1981; Hunzelmann and Spillmann, 1984; Schieting and Spillmann, 1987; Anstis, 1996). It merely means that statistically they are less likely to disappear. It is like the old joke about the married couple being pursued by a hungry grizzly bear. As the wife runs she (yellow spots) reflects that she need not outrun the bear (blue spots), she need only outrun her husband. Specifically, the moving blue dots resist disappearance more than the yellow dots, but they do not *induce* any blindness to the blue dots. The situation is like a horse race in which the winner does not slow down the losers, as opposed to a tug of war in which the winners do oppose and impede the efforts of the loser. The wife does not need to trip up her husband, only outrun him. (Incidentally, eye movements will shift all the spots across the retina and tend to keep them visible.)
3. *The disappearances are generally not noticed unless attention is specifically directed toward them.*

I call this process "fade blindness." In this model, *all* the spots disappear from time to time. The static yellow spots may disappear more often than the moving blue spots, but in addition only the yellow ones are salient enough, and receive enough attention, for us to notice their fading. Consider any peripherally viewed blue spot. If this were to fade, its fading would scarcely be noticed, since it is already randomly flickering and changing luminance and is harder to distinguish from the other blue spots in any case. However, the three yellow spots stand out very sharply from the surround, and if these fade, it is immediately noticed. I suggest that *all* the spots stand an almost equal chance of fading, but the observer only notices this when the spots happen to be yellow and hence pop out and seize his attention. Thus all spots are constantly subjectively disappearing and reappearing, but only those spots that capture the observer's attention are correctly seen to be fading. To demonstrate this, three flickering blue dots were painted magenta and three others were painted green. The nine dots (yellow, magenta, and green) were arranged in an irregular circle around the fixation point. Result: the yellow spots still showed dramatic disappearances. But when observers were asked to concentrate on the magenta spots, or on the green spots, these attended spots also disappeared at irregular times. The purpose of the magenta and green colors was simply to label some of the flickering dots to make it easier for the observer to attend to them. The fact that these disappeared supports our conclusion that any dot can fade out, and hence we conjecture that all of them do so at irregular intervals. The fading escapes our notice, rather as the temporal changes in change-blindness movies also escape our notice, until we specifically direct our attention toward them (see below).

This hypothesis predicts that selectively paying attention to some peripheral spots will paradoxically make them more likely, not less likely, to fade out. Lou (1999) reported exactly this effect. He presented a circular array of six disks — three green and three orange — in alternate positions, against a uniform gray background. Sixteen observers maintained steady fixation at the center of the array and were instructed to direct their attention to three disks of one color and to ignore the three disks of the other color. In about 10 s, some disks started to fade away from awareness. Of those starting to fade, 81% were those selected for attention. The faded disks remained out of awareness for 1–2 s during which time other disks were clearly visible. The fading increased with eccentricity, a defining characteristic of Troxler fading. Lou concluded from the selectivity of the fading that voluntary attention can have an inhibitory effect on early sensory processing. However, we take a rather different view; we suggest that attention did not alter the sensory processing in any way, but simply made the Ss aware of the fading that they (and indeed all of us) usually ignore.

It may seem bizarre to suggest that all our lives peripheral stimuli have been fading without our ever being aware of it. But an analogous situation has recently been discovered, in which our picture of the world is far more fragile and impoverished than everyday intuition would suggest. I refer to change blindness. Suppose that a photograph of a jet plane is flashed up, followed immediately by a doctored version of the same scene that has been changed in some obvious way, say by removal of one jet engine, and these two pictures cycle continuously. Observers immediately report the location and nature of the change. However, if this scenario is repeated with fresh observers, but now with a blank gray interstimulus interval (ISI) of 250 ms inserted after each picture (A-blank-B-blank-A-blank…), observers typically fail to notice the change and need to scan the alternating picture sequence carefully for up to 30 s before they notice the change. The conclusion is that in the first instance, the engine flickering on and off captures observer's attention and this cues the identification of the change. However, inserting the ISI generates flicker transients over the whole field. Without the local attention-grabbing transient, the observer fails to see the change. This remarkable "change blindness" forces us to abandon the traditional belief that the visual system builds up a detailed and complex picture of the world over time, and suggests instead that our visual representations are sparse and volatile (Dennett, 1991; O'Regan, 1992; Simons and Levin, 1997; Wolfe, 1998; O'Regan et al., 1999; Rensink, 2000; Simons, 2000). Careful parametric studies by Becker and colleagues (Becker et al., 2000; Becker and Pashler, 2002) confirm that our visual inputs are indeed far more impoverished than we usually imagine. I am conjecturing that "fade blindness" is a near-universal but hitherto unrecognized phenomenon that, together with change blindness, seriously restricts the amount of information that we actually take in from the world. Careful attentive scrutiny is necessary to detect both the vanishing jet engine and the vanishing peripheral spots.

Adapting to irregularity

Some of the peripherally viewed stimuli listed earlier actually disappear during strict fixation. Others remain visible but they lose some of their visual properties; for instance, moving or flickering objects do not vanish but they gradually seem to lose their motion (Campbell and Maffei, 1979; Hunzelmann and Spillmann, 1984) or flicker (Schieting and Spillmann, 1987; Anstis, 1996). In such cases the visual system is adapting to higher order properties such as motion, flicker, or texture (e.g., Anstis, 1983), rather than to the luminance contours that define the object's existence and location. Alan Ho and I have measured visual adaptation to an unusual and frequently overlooked visual property, namely geometrical *irregularity*. We found that peripherally viewed wiggly or curved lines, or irregularly arranged dots, remain visible but gradually look smoother and more regular. This hints at the way in which the visual system codes for spacing and curvature.

Figs. 6a, b show a set of wiggly vertical lines, wiggly circles, and irregularly arranged dots, each mirrored about a fixation point (Mackay, 1964a, b). In each case, cover the right-hand half with a piece of paper and gaze steadily at the fixation point for 30 s or so. Now remove the paper and you will notice that the adapted lines, circles, or dots on the left look much straighter and more regular than the freshly exposed ones on the right. Note that in all these cases the test and the adapting stimulus were identical — the stimulus changed its appearance during prolonged inspection. This is different from the commoner measure of an aftereffect in which (say) a tilted adapting line alters the appearance of a vertical test line. Our process was a form of normalization, comparable to the way in which a tilted line gradually normalizes toward the vertical (Gibson and Radner, 1937).

Wiggly lines. In our experiments, a fixed set of adapting wiggly vertical lines lay to the left of a fixation point. Once per second an adjustable set of wiggly test lines was flashed up to the right of the fixation point for 100 ms. These test lines were a mirror image of the adapting lines, except that their amplitude was under the observer's control, and she/he hit keys to adjust this amplitude in

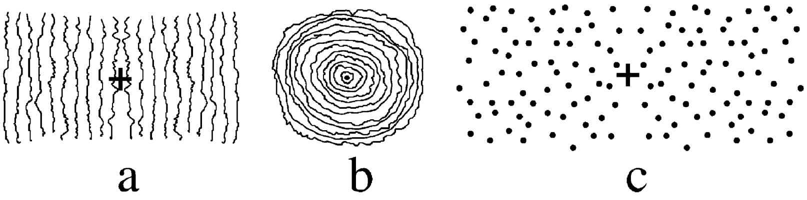

Fig. 6. (a)–(c) Adapt to irregularity (Mackay, 1964a, b). Cover one half, adapt for 30–60 s. Then uncover the half that should look more irregular.

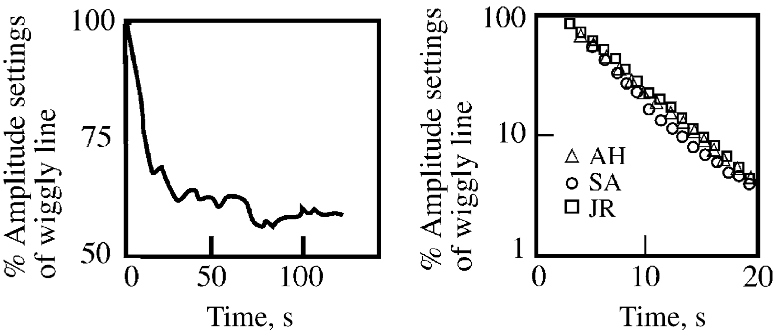

Fig. 7. (a) Illusory smoothing of wiggly lines over time. The y-axis shows the amplitude or gain of a matching wiggly line that subjectively matched a constant wiggly test line over time. (b) Illusory smoothing of sinusoidal curved lines over log time, averaged across spatial frequencies. All three observers had a time constant of 2.5 s. Note different scales on y-axis.

order to match subjectively the gradual decline of wiggliness in the adapting lines. Result: the illusory smoothing over time is shown in Fig. 7a.

Sinusoidal curves. Mathematically, every wiggly line is composed of a set of sinusoids of different amplitudes and spatial frequencies. So we also adapted to a single sinusoidal line, flashing up an adjustable comparison sinusoid next to it. Result: Fig. 7b shows that averaging across a range of different spatial frequencies, the perceived amplitude fell off exponentially with time, with a time constant of 2.5 s, the same for all three observers.

Irregular dot lattice. First, an irregular lattice was generated from a regular square lattice by sending each dot on a random walk, like a diffusion process. This adapting lattice was steadily fixated, and an adjacent, adjustable matching lattice was flashed up once every second for 100 ms, and again the observer continuously adjusted its irregularity to match the gradually changing appearance of the adapting irregular lattice. Result (not shown): the perceived regularity increased monotonically over time.

Interpretation. Instead of suggesting that these irregular stimuli became gradually straighter or more regular, I claim that they became less irregular. In other words, irregularity is a class of visual property that is explicitly coded by the visual system, analogous to color. Just as there are different hues, so there are different forms of irregularity, the most obvious being variance in curvature or in spacing. And just as adapting to colors tends to reduce their saturation, adaptation reduces the strength of irregularity. In support of this idea, we have also found that sinusoidally curved lines looked less curved (i.e., straighter) when they were masked with twinkling dynamic noise, or when they were defined by second-order texture or by twinkle instead of by luminance. We conclude that masking noise simulates the effects of prolonged adaptation, because both reduce the neural signal/noise ratio, in this case whatever it is that signals curvature. This was especially true of low-amplitude sinusoids. Thus, noise masking and prolonged adaptation both serve to reduce the signal/noise ratio of the neural signals of irregularity. So irregularity is a signal; added to a straight line, it gives curved or wiggly lines; added to a regular square lattice, it transforms it into an irregular one.

Visual search and added properties. Treisman and Souther (1985) showed that it is easier to find a Q hidden among O's than it is to find an O hidden among Q's. They argued that the tail of Q is an added visual feature or property, and it is easier to find a target that possesses the feature among distracters that lack it than vice versa. So if curvature is an added property, then curved lines should be easy to find among straight lines, whereas straight lines should be hard to find among curved lines. And they are (Treisman and Gormican, 1988).

<table>
<tr><th>Easy</th><th>Hard</th></tr>
<tr><td>OOQOO</td><td>QQOQQ</td></tr>
<tr><td>|||)|||</td><td>)))|)))</td></tr>
</table>

When searching arrays of dot lattices, it should be easier to find an irregular patch of dots hidden among regular dots than vice versa. Sturzel and Spillmann (2001) adapted to patches of irregular dots hidden among regular lattices, and found that they vanished after some seconds. This is similar to the adaptation effects that we found for lattices.

A differential-adaptation model. Prolonged inspection of a sinusoidal line adapts out the curvature signal so that the line gradually comes to look straighter. We conjecture that a single long receptive field running along the line codes its overall (mean) orientation, while the local curves are coded by small, end-stopped receptive fields parallel to the local curves. If these smaller units adapt out more rapidly than the larger long unit, the line will gradually come to look straighter.

Again, a tilted line is probably coded by overlapping, differently oriented receptive fields. Suppose that a line tilted 10° away from the vertical is coded by equal firing from one unit tuned to 0° (vertical) and a second unit tuned to 20° away from vertical. If the second more tilted unit adapts more rapidly, the tilted line will appear to become gradually more vertical. In general, we can model the normalization of any "self-adapting" stimulus by a firing-ratio model consisting of one unit tuned to a "norm," such as straight or vertical, and a second unit tuned somewhere away from the norm, such as oriented at some angle away from the vertical. If the norm unit adapts more slowly than the other unit, the stimulus will appear to normalize over time.

Glaucoma prediction. This leads to a clinical prediction. It has been shown that chronic glaucoma selectively damages large optic nerve fibers, which generally have large receptive fields. The loss of large cells tends to be greater in the lower than the upper retina, and to spare the fovea (Quigley et al., 1987, 1988; Glovinsky et al., 1991). Our model predicts that a sinusoidally curved line will look *more curved* to a glaucomatous patient because the straight-line signal from the long receptive field will be attenuated compared to the local curves signaled by the still intact smaller receptive fields. Projecting two identical sinusoidal lines on to more and less seriously affected retinal regions, such as the lower and upper retina of the same patient, will allow the observer to act as his own control; the more extensive the large cell loss in any retinal region, the more curved the line should look.

It might be possible to mimic the supposed visual effects of selective loss of large or small cells by adapting a normal observer to isotropic (nonoriented) dynamic noise that had been spatially filtered to contain only high or only low spatial frequencies. Preadaptation to coarse dynamic noise should temporarily attenuate large receptive fields and briefly simulate the visual experience of glaucoma. It should make a sinusoidally curved line look more curved. Adaptation to fine noise should have the opposite effect and make the line look apparently straighter.

A failure: do correlated sensory inputs teach the brain where things are?

Most of the experiments I do are rather dull, with results that would not surprise anybody. But at least they are very likely to work and lead to publications that keep my merit committee happy. But now and again it is worth attempting an experiment with a low probability of success but a very high payoff if it does succeed. I shall now describe such an attempt. The low probability won out over the high payoff and the experiment did not succeed, to put it mildly, but I am still glad that I tried it. The idea was to see whether the visual system can modify or rewire itself in response to experimentally correlated visual inputs.

Crick and Koch (2003, p. 120) asked:

> How are feature detectors formed? A broad answer is that neurons do this by detecting common and significant correlations in their inputs and by altering their synapses (and perhaps other properties) so that they can more easily respond to such inputs. In other words, the brain is very good at detecting apparent causation. Exactly how it does this is more controversial. The main mechanism is probably Hebbian, but Hebb's seminal suggestion needs to be expanded.

I narrowed down Crick and Koch's question to ask: how do we know where things are in space? Are we born with an innate knowledge of retinal sign, or do we acquire this by some learning process? One can imagine the brain to be like a TV technician sitting in a studio and receiving visual signals that arrive up the million cables of the optic nerve. At first he has no idea as to which signal comes from where in the retina, but by careful inspection he notices that some cables carry correlated signals. For instance, a pair of vertically aligned retinal receptors is likely to fire together when a vertical edge passes over them. Adjacent receptors are statistically likely to see the same colors, luminances, edge orientations, directions of motion, and so on. When he notices two cables with correlated inputs he ties them together. Gradually he can arrange the cables in a two-dimensional pattern that corresponds to the layout of the retina. Maloney and Ahumada (1989) offer some hints on how this might be done.

Much evidence from brain scans and neural recordings indicates that the brain may be an expert statistician that can draw inferences from correlated sensory inputs via Hebbian learning (reviewed by Cruikshank and Weinberger, 1996). Consider first the somatosensory system. Each body region sends signals to its own brain area. For instance, the five fingers on one hand all send signals to adjacent patches of the somatosensory cortex. Allard et al. (1991) glued together two fingers of an owl monkey's hand, and found that the two brain areas of the two fingers eventually coalesced into a single brain area. Were these brain changes produced by correlated inputs or correlated outputs? Gluing the two fingers together would correlate the motor outputs, since the two fingers were forced to move together, but it would also correlate the sensory inputs, since the two fingers would nearly always touch the same surfaces at the same time. Wang et al. (1995) showed that correlated inputs produced the brain changes. Without gluing the fingers together, they made a mechanical tapper that tapped both fingers simultaneously all day long at random times. Again the two brain areas coalesced. It occurred to me to do the same thing in vision, by applying correlated visual inputs to adjacent retinal areas.

The visual cortex also changes its firing patterns in response to correlated visual inputs (e.g., Eysel et al., 1998). Some computational models have suggested how the brain might wire up ''local sign'' on the retina by means of an unsupervised learning algorithm, based upon correlations of the retinal inputs that result from eye movements (Maloney and Ahumada, 1989).

Vision is obviously a more highly developed sense than skin sensitivity, so it seemed like a good idea to investigate correlated inputs in vision rather than on the skin. A prime example of correlated visual inputs is stereo vision, where the two eyes receive almost identical pictures which the brain can fuse or combine, and can use the small differences between the two pictures to calculate depth. Is this process hard wired or learned? There is much evidence that the two eyes compete for ''brain space'' during early development. Whereas acuity develops gradually over the first 6 months of life, stereo vision appears suddenly, usually within a few days, during the 20th week of life (reviewed by Atkinson, 2000). This suggests that stereo vision is hard wired, though it could in principle involve latent learning that shows no signs until a sudden winner-take-all decision process permits binocular fusion.

I correlated the visual inputs, not by putting near-identical pictures one into each eye but by putting both pictures side by side into the same eye(s), stimulating two nearby retinal areas with the same regime. Originally I thought of using two light-emitting diodes that would flash and flicker in step at the same rate, with matching brightness and color. But a much better idea is to expose two identical movies side by side on a split-screen television. With some difficulty I persuaded my university to buy me a new flat-panel TV with a built-in split-screen facility, which I installed in my living room. Whatever picture was exposed on the TV screen, it appeared as two identical pictures side by side and touching. Each picture was 7° wide at the viewing distance used, so that points that were horizontally separated spatially by 7° (one screen width) were always identical in color, flicker, motion, and brightness. Thus, the split-screen television set had the identical program on the two halves of the screen. Observers were instructed to watch the left-hand picture. There was now a point-to-point matching of brightness, color, edges, flicker, and movement over a considerable retinal area. Note that extremely strict fixation was not necessary; provided that the observer always looked somewhere in the middle of the left-hand picture, there was a large region ($<7°$) where points exactly 7° apart were correlated (see Fig. 8).

I invited three student research assistants over for the day, and all of us watched TV all day for six consecutive hours. We found that the correlation between the two screens felt “weird.” This correlation was most obvious during periods of maximum motion and activity in the picture, and also whenever cuts between scenes, or camera zooms and pans, produced sudden massive changes in the picture. Of course the two pictures were always correlated 100%, but this was subjectively far more obvious for moving than for static pictures.

From time to time we ran some informal visual tests, looking for small perceptible changes. For instance, we thought it possible that correlated points would gradually begin to look closer together, so we ran a bisection task in which three dots were horizontally separated by $\sim 7°$. Observers adjusted the positions of the dots until the middle dot appeared to lie halfway between the two outer dots. We predicted that observers might underestimate the distance between the middle and the right-hand dots, corresponding to the separation between the fixated and the peripherally viewed screen.

We also thought that the visual system might start to predict the existence of the correlations, which might be visible as a form of contingent aftereffect (Shute, 1979). Accordingly, we occasionally covered the peripherally viewed screen with a black card to see whether observers could notice any ghostly duplicate of the centrally viewed picture. Also, we sometimes turned off the pair of pictures and substituted a single, foveally viewed stimulus such as a red cross rotating clockwise. Observers looked for any hint of an illusory duplicate shape 7° to the right, which might take the form of a ghostly red or green cross rotating clockwise or counterclockwise.

Unfortunately, all these tests came out negative. We never saw the slightest errors in bisection, nor

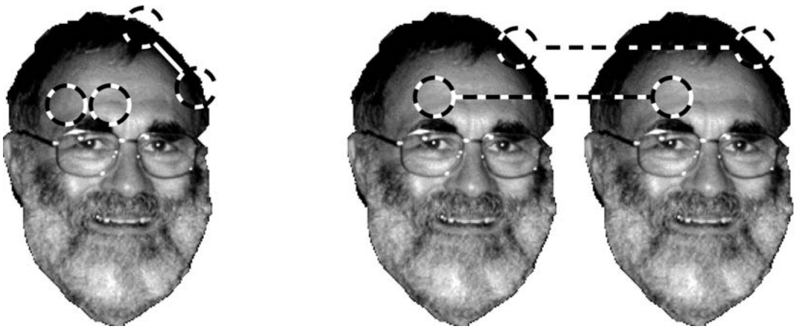

Fig. 8. (a) In normal vision, nearby points are statistically alike in orientation (hairline), color, and texture (forehead) ($0<p<1$). (b) In split-screen mode, points separated by 7° are precisely matched in orientation *and* color *and* texture and so on ($p = 1$).

ghostly illusions nor visual disturbances of any kind. So the experiment was a failure. However, it still looks like a promising line of enquiry, and given a better-designed experiment or a smarter experimenter it might still pay off. Perhaps I can null out the effects of eye movements if I spend more money.... My only problem is that my department may soon demand to get its TV back. When the TV goes, I shall have plenty of time to re-read Lothar Spillmann's huge collected works.

Conclusions

Among the new findings summarized here are the following:

1. Neon spreading is strongest at equiluminance. It does not necessarily require intersecting continuous lines, but can be seen even on scattered random dots.
2. The "dam" theory can explain the fading out of peripheral stimuli during adaptation, but not the subsequent aftereffects that are often misnamed as "filling-in."
3. We propose a speculative theory to explain phenomena of peripheral fading, in particular motion-induced blindness, which we suggest may not really be motion induced. In brief, all peripheral stimuli tend to disappear, especially during fixation, so we have a much sketchier picture of the peripheral scene than we generally believe. It is only when we attend to specific peripheral items that we become aware that they are fading. We regard this as a new cousin of change blindness.
4. During steady fixation, irregular lines, grids, and scattered dots gradually look more regular, or to be more accurate less irregular. We claim that irregularity or randomness is explicitly coded as a visual dimension, and adaptation reduces the strength of this coded neural signal.

References

Allard, T., Clark, S.A., Jenkins, W.M. and Merzenich, M.M. (1991) Reorganization of somatosensory area 3b representations in adult owl monkeys after digital syndactyly. J. Neurophysiol., 66: 1048–1058.

Anstis, S.M. (1979) Interactions between simultaneous contrast and adaptation to gradual change of luminance. Perception, 8: 487–495.

Anstis, S. (1983) Aftereffects of form, motion, and color. In: Spillmann, L. and Wooten, B.R. (Eds.), Sensory Experience, Adaptation, and Perception. Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 583–601.

Anstis, S. (1989) Kinetic edges become displaced, segregated and invisible. In: Lam, D.M.K. and Gilbert, C.D. (Eds.), Neural Mechanisms of Visual Perception. Gulf Publishing Company, Houston, pp. 247–260.

Anstis, S. (1996) Adaptation to peripheral flicker. Vision Res., 36: 3479–3485.

Anstis, S.M. and Reinhardt-Rutland, A.H. (1976) Interactions between motion aftereffects and induced movement. Vision Res., 16: 1391–1394.

Anstis, S., Rogers, B. and Henry, J. (1978) Interactions between simultaneous contrast and coloured afterimages. Vision Res., 18: 899–911.

Atkinson, J. (2000) The Developing Visual Brain. Oxford University Press, Oxford.

Baumgartner, G. (1960) Indirekte Größenbestimmung der rezeptiven Felder der Retina beim Menschen mittels der Hermannschen Gittertäuschung. Pflügers Arch. gesamate Physiol., 272: 21–22.

Becker, M.W. and Pashler, H. (2002) Volatile visual representations: failing to detect changes in recently processed information. Psychonom. Bull. Rev., 9: 744–750.

Becker, M.W., Pashler, H. and Anstis, S. (2000) The role of iconic memory in change-detection tasks. Perception, 29: 273–286.

Bonneh, Y.S., Cooperman, A. and Sagi, D. (2001) Motion-induced blindness in normal observers. Nature, 411: 798–801.

Bressan, P., Mingolla, E., Spillmann, L. and Watanabe, T. (1997) Neon color spreading: a review. Perception, 26: 1353–1366.

Campbell, F.W. and Maffei, L. (1979) Stopped visual motion. Nature, 278: 192.

Crick, F. and Koch, C. (2003) A framework for consciousness. Nat. Neurosci., 6(2): 119–126.

Cruikshank, S.J. and Weinberger, N.M. (1996) Evidence for the Hebbian hypothesis in experience-dependent physiological plasticity of neocortex: a critical review. Brain Res. Rev., 22: 191–228.

Daw, N.W. (1962) Why after-images are not seen in normal circumstances. Nature, 196: 1143–1145.

De Weerd, P., Desimone, R. and Ungerleider, L.G. (1998) Perceptual filling-in: A parametric study. Vision Res., 38: 2721–2734.

Dennett, D.C. (1991) Consciousness Explained. Little, Brown and Co., Boston.

Eysel, U.T., Eyding, C.A.D. and Schweigart, G. (1998) Repetitive optical stimulation elicits fast receptive field changes in mature visual cortex. Neuroreport, 9: 949–954.

Geier, J., Sera, L. and Bernath, L. (2004) Stopping the Hermann grid illusion by simple sine distortion. Perception ECVP Abstracts.

Gibson, J.J. and Radner, M. (1937) Adaptation, after-effect, and contrast in the perception of tilted lines. I. Quantitative studies. J. Exp. Psychol., 20: 453–467.

Glovinsky, Y., Quigley, H.A. and Dunkelberger, G.R. (1991) Retinal ganglion cell loss is size dependent in experimental glaucoma. Invest. Ophthalmol. Vis. Sci., 32: 484–491.

Grassman, H. (1853) Zur Theorie der Farbenmischung. Poggendorfs Ann. Phys., 89: 69–84.

Hermann, L. (1870) Eine Erscheinung simultanen Contrastes. Pflügers Arch. gesamte Physiol., 3: 13–15.

Hunzelmann, N. and Spillmann, L. (1984) Movement adaptation in the peripheral retina. Vision Res., 24: 1765–1769.

Kelly, D.H. (1983) Spatiotemporal variation of chromatic and achromatic contrast thresholds. J. Opt. Soc. Am., 73: 742–750.

Kitaoka, A., Gyoba, J., Kawabata, H. and Sakurai, K. (2001) Two competing mechanisms underlying neon color spreading, visual phantoms and grating induction. Vision Res., 41: 2347–2354.

Lou, L. (1999) Selective peripheral fading: evidence for inhibitory sensory effect of attention. Perception, 28: 519–526.

Mackay, D.M. (1964a) Central adaptation in mechanisms of form vision. Nature, 203: 992–993.

Mackay, D.M. (1964b) Dynamic distortions of perceived form. Nature, 203: 1097.

Maloney, L.T. and Ahumada, A.J. (1989) Learning by assertion: two methods for calibrating a linear visual system. Neural Comput., 1: 392–401.

Ninio, J. and Stevens, K.A. (2000) Variations on the Hermann grid: an extinction illusion. Perception, 29: 1209–1217.

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

O' Regan, K.J., Rensink, R.A. and Clark, J.J. (1999) Change blindness as a result of mudsplashes. Nature, 398: 34.

Oehler, R. and Spillmann, L. (1981) Illusory colour changes in Hermann grids varying only in hue. Vision Res., 21: 527–541.

Pessoa, L. and deWeerd, P. (2003) Filling-in: From Perceptual Completion to Cortical Reorganization. Oxford University Press, Oxford.

Pinna, B., Brelstaff, G. and Spillmann, L. (2001) Surface color from boundaries: a new 'watercolor' illlusion. Vision Res., 41: 2669–2676.

Pinna, B., Werner, J.S. and Spillman, L. (2003) The watercolor effect: a new principle of grouping and figure-ground organization. Vision Res., 43: 43–52.

Quigley, H.A., Dunkelberger, G.R. and Green, W.R. (1988) Chronic human glaucoma causing selectively greater loss of large optic nerve fibers. Ophthalmology, 95: 357–363.

Quigley, H.A., Sanchez, R.M., Dunkelberger, G.R., L'Hernault, N.L. and Baginski, T.A. (1987) Chronic glaucoma selectively damages large optic nerve fibers. Invest. Ophthalmol. Vis. Sci., 28: 913–920.

Ramachandran, V.S. and Gregory, R.L. (1991) Perceptual filling in of artificially induced scotomas in human vision. Nature, 350: 699–702.

Ransom-Hogg, A. and Spillmann, L. (1980) Perceptive field size in fovea and periphery of the light- and dark-adapted retina. Vision Res., 20: 221–228.

Redies, C. and Spillmann, L. (1981) The neon color effect in the Ehrenstein illusion. Perception, 10: 667–681.

Redies, C., Spillmann, L. and Kunz, K. (1984) Colored neon flanks and line gap enhancement. Vision Res., 24: 1301–1309.

Rensink, R.A. (2000) The dynamic representation of scenes. Vis. Cogn., 7: 17–42.

Schieting, S. and Spillmann, L. (1987) Flicker adaptation in the peripheral retina. Vision Res., 27: 277–284.

Schrauf, M., Lingelbach, B. and Wist, E.R. (1997) The scintillating grid illusion. Vision Res., 37: 1033–1038.

Schrauf, M. and Spillmann, L. (2000) The scintillating grid illusion in stereo-depth. Vision Res., 40: 717–721.

Shute, C.C.D. (1979) The McCollough effect. Cambridge University Press, Cambridge.

Simons, D.J. (2000) Current approaches to change blindness. Vis. Cogn., 7: 1–16.

Simons, D.J. and Levin, D.T. (1997) Change blindness. Trends Cogn. Sci., 1: 261–267.

Spillmann, L. (1971) Foveal perceptive fields in the human visual system measured with simultaneous contrast in grids and bars. Pflugers Archiv., 326: 281–299.

Spillmann, L. (1994) The Hermann grid illusion: a tool for studying human perspective field organization. Perception, 23: 691–708.

Spillmann, L. (1999) From elements to perception: local and global processing in visual neurons. Perception, 28: 1461–1492.

Spillmann, L. (2003) Re-viewing 25 years of ECVP — a personal view. Perception, 32: 777–791.

Spillmann, L. and de Weerd, P. (2003) Mechanisms of surface completion: perceptual filling-in of texture. In: Pessoa, L. and deWeerd, P. (Eds.), Filling-in: From Perceptual Completion to Cortical Reorganization. Oxford University Press, Oxford, pp. 81–105.

Spillmann, L. and Kurtenbach, A. (1992) Dynamic noise backgrounds facilitate target fading. Vision Res., 32: 1941–1946.

Spillmann, L. and Levine, J. (1971) Contrast enhancement in a Hermann grid with variable figure-ground ratio. Exp. Brain Res., 13: 547–559.

Spillmann, L. and Redies, C. (1981) Random-dot motion displaces Ehrenstein illusion. Perception, 10: 411–415.

Spillmann, L. and Werner, J.S. (1996) Long-range interactions in visual perception. Trends Neurosci., 19: 428–434.

Sturzel, F. and Spillmann, L. (2001) Texture fading correlates with stimulus salience. Vision Res., 41: 2969–2977.

Treisman, A. and Gormican, S. (1988) Feature analysis in early vision: Evidence from search asymmetries. Psychol. Rev., 95: 15–48.

Treisman, A. and Souther, J. (1985) Search asymmetry: a diagnostic for preattentive processing of separable features. J. Exp. Psychol. Gen., 114: 285–310.

Troxler, D. (1804) Ueber das Verschwinden gegebener Gegenstaende innerhalb unseres Gesichtskreises. In: Himly, K. and Schmidt, J.S. (Eds.), Ophthal. Bibliothek II.2. F Frommann, Jena, pp. 1–119.

Tyler, C.W. (1973) Temporal characteristics in apparent movement: omega movement vs. phi movement. Quart. J. Exp. Psychol., 25: 182–192.

Tyler, C.W. and Hardage, L. (1998) Long-range twinkle induction: An achromatic rebound effect in the magnocellular processing system. Perception, 27: 203–214.

Wang, X., Merzenich, M.M., Sameshima, K. and Jenkins, W.M. (1995) Remodelling of hand representation in adult cortex determined by timing of tactile stimulation. Nature, 378: 71–75.

Wolfe, J.M. (1984) Global factors in the Hermann grid illusion. Perception, 13: 33–40.

Wolfe, J.M. (1998) Visual memory: What do you know about what you saw. Curr. Biol., 8: R303–R304.

Zaidi, Q. (1999) Color and brightness induction: from Mach bands to three-dimensional configurations. In: Boynton, B.B., Gegenfurtner, K.R. and Sharpe, L.T. (Eds.), Color Vision: From Genes to Perception. Cambridge University Press, Cambridge, pp. 317–343.