Multiple Analyses of Orientation in the Visual System

Patrick Cavanagh

Département de Psychologie
Université de Montréal
Montréal, Québec
H3C 3J7

We have examined aftereffects and illusions using images defined by five different visual attributes: luminance, color, texture, motion, and binocular disparity. Our results suggest that, for all five attributes, the initial representation of two-dimensional shape involves information about local size and orientation and that in many cases, these size and orientation codes are analyzed independently for individual attributes. Why should these two-dimensional size and orientation analyses be duplicated for several attributes? First, multiple analyses increase reliability. In particular, color and texture borders are more reliably linked to object borders than is luminance (luminance contours are often confounded by shadow borders). Second, using similar codes for each attribute allows a standardized image description that facilitates the subsequent integration of the separate images.

Why are the low-level codes specifically size and orientation? Again reliability may be the reason. Early contour extraction is significantly more accurate when based on local size and orientation information, and the ultimate role of these codes may be to develop a contour representation of the image. On the other hand, difficulties in classifying different types of image contours suggest that the initial memory access should not be based on explicit contour representations but on representations that are invariant to size and orientation. These also depend on the early extraction of size and orientation information. The functionally independent pathways that are demonstrated in the experiments reported here do not appear to do much more than extract two-dimensional representations. It seems very unlikely that such an early stage in visual processing would involve all the visual cortices up to and including areas MT and V4. It is proposed that these cortices are the beginnings of a different set of pathways that are not specific to stimulus attributes.

Introduction

The goal of vision is to inform us both of the identity of objects in view and their spatial positions. Research in several fields has recently indicated that the initial information concerning the scene may be analyzed as several separate images, each involving a particular visual attribute such as texture, color, or motion. Each of these separate images conveys a message about the value of the attribute in question, indicating whether a particular position is, say, red or green. Some of these images also provide three-dimensional information such depth from binocular disparity or surface slants from local gradients of motion or texture. Although the analysis of these attribute-specific

images likely involves specialized processes, all of the images convey a message about the two-dimensional shape of regions, and all of these may represent two-dimensional shape in a similar manner. These two-dimensional shape descriptions can then be used to infer three-dimensional object structure through cues such as occlusion, perspective, contour junctions, and the change of shape over time (structure from motion). The extraction of a two-dimensional description of the retinal image is a major step in visual analysis and the purpose of this paper is to examine the nature and degree of independence of the analyses of two-dimensional shape for different visual attributes.

Physiological studies have given clear evidence of multiple analyses at early levels in vision. An image can be considered a superposition of several surface attributes such as color, texture, and luminance and in the first area of visual cortex, V1, cells respond to many of these attributes conjointly: orientation, size, color, direction of motion, binocular disparity. We could imagine processes that operate on this combined image in order to segment regions defined by various attributes but that does not seem to be the case. Cowey[1] and Barlow[2] have argued that local interactions are an integral part of the low-level processes that the visual system applies to these attributes and that for reasons of simplicity and economy (length of axons), the visual system has opted to analyze different attributes in physically separate areas. There is physiological evidence that

relative motion and color are analyzed by specialized cortical areas[3,4] and therefore that the stimulus information defined only by color or only by motion will follow separate pathways through these areas. Visual deficits following brain lesions in human patients have shown independent losses of vision for motion,[5] color,[6] and luminance,[7] indicating some physical separation in the cortical representations or projections for these attributes. Although the existence of multiple representations in extrastriate cortex is widely accepted, the routing of information through these areas and the stimulus and/or response attributes processed by each are not yet clearly established.[8]

The work described in this paper examines the functional pathways of information by using perceptual tasks in normal humans. We have studied five stimulus attributes (Figure 1): color, luminance, texture, binocular disparity, and relative motion. Undoubtedly, the actual pathways are considerably more tortuous than those shown in Figure 1, but this simplified schematic can serve as a useful point of departure for research. To construct images defined by a single attribute, texture for example, we start with a black and white figure and replace the black areas with one texture and the white areas with a different texture having the same mean luminance. Figure 2 demonstrates this attribute replacement for texture and binocular disparity. Using images defined by individual attributes, we have been able to demonstrate psychophysically that there are independent analyses of several attributes.[9-12] Whether or
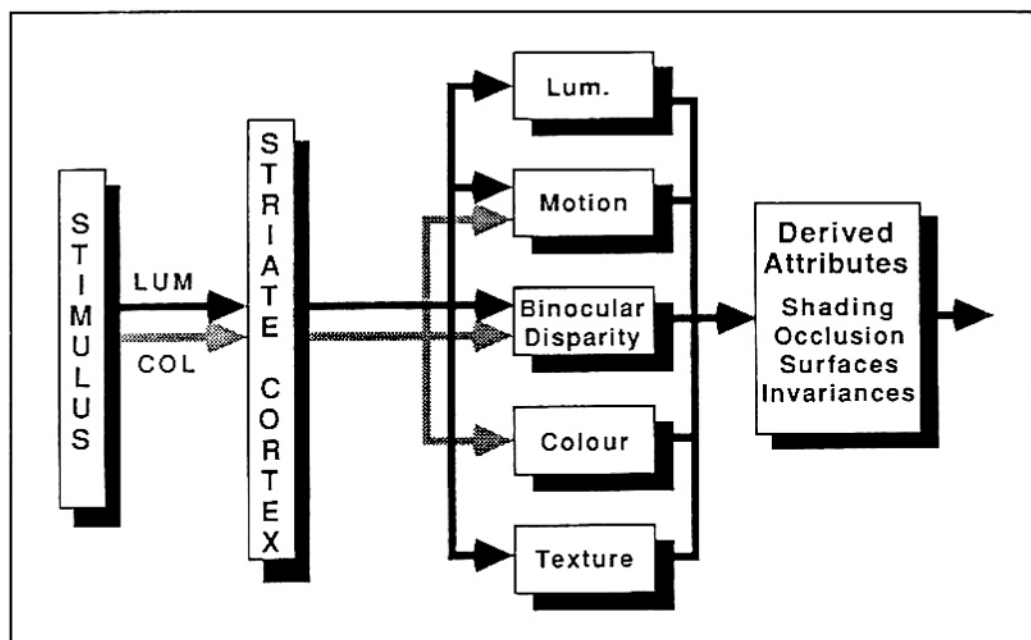
Figure 1. Perceptual pathways in the visual system. Luminance, motion, binocular disparity, color, and texture are stimulus surface qualities that may receive independent analyses. Each of these analyses generates a two-dimensional representation of the attribute, contributing to an overall representation of stimulus shape. These independent representations may be followed by one or more high-level representations that reintegrate information from all attributes and from which are drawn inferences of shading, occlusion, and surfaces inferences.

not analyses that are functionally independent are also physically separate is a question best answered with physiological methods. In fact, since several studies have shown that analysis of different attributes can be compartmentalized within one visual cortex (for example, processing of chromatic properties in the cytochrome oxidase blob areas of V1 and achromatic properties in interblob regions),[13,14] functional independence may imply either physical separation in local compartments or in separate cortices.

Whatever the routing of information through various stages of early analysis, the separate descriptions must at some point be recombined, and this raises interesting questions concerning

the extent of the independent analysis, the cooperation across attributes, and the resolution of conflicts. Bulthoff and Mallot[15] and Sperling and Dosher[16] have described situations of conflict between various cues in establishing three-dimensional shape. Gregory[17] has proposed that luminance is the primary attribute in determining contour location and that other attributes simply fill out to the luminance-defined border. He calls this border locking and claims that the luminance information is the master map and that the others are adjusted into register with it. Our work[9-11,18-22] gives no evidence of special primacy or privilege for luminance information other than the extra resolution it affords. We feel that the determi-
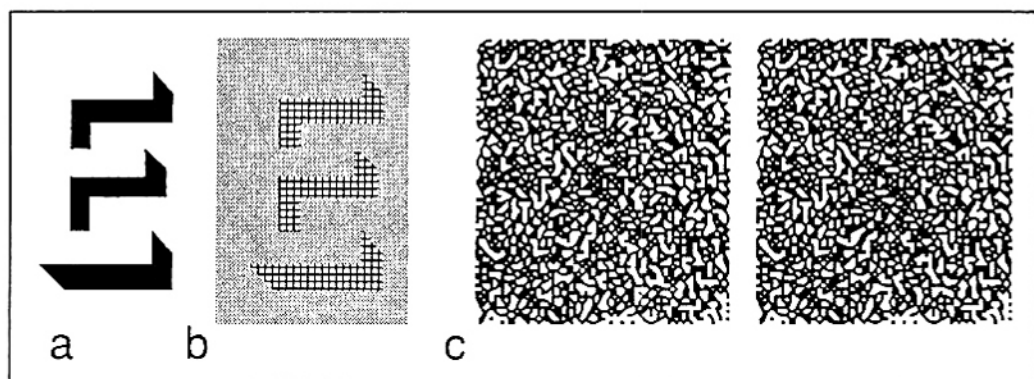
Figure 2. Representations of a shadowed, block letter E defined by (a) luminance, (b) texture, and (c) binocular disparity (the left and right hand panels must be fused).

nation of object borders involves much more interattribute cooperation. In fact, the redundancy offered by the independent analysis of separate attributes may be the most important advantage of separate analyses.

What follows this initial two-dimensional analysis of visual attributes? Are rudimentary object descriptions built up from these two-dimensional images of individual attributes at an early level in vision? The object-based approaches of Marr[23] and Biederman[24] suggest that they are. Work that we have just begun[25,26] takes a quite different approach, suggesting that object parts and boundaries should not be explicitly identified at such an early stage and that matching of raw two-dimensional views may be the most effective way to make the initial memory contact. I will not describe the nature of object representations here but concentrate rather on the nature of the multiple representations in early vision.

Inputs. Although Figure 1 places all the five attributes we have considered at the same level, the different attributes, in fact, become explicitly represented at different stages in the visual system. As shown schematically in Figure 3, luminance and color are the initial attributes that represent the image, and these are followed by binocular disparity, motion, and texture, each of which emerges as a property of contours defined by the first two.

Since there is an area specialized for the analysis of color (V4, 3) it is often assumed that color information does not contribute to the pathways involved in the analysis of motion and binocular disparity. However, color does contribute to motion[27-30] and to binocular disparity.[31,32] Our most recent studies[30,33] suggest that the contribution of color to motion passes through the opponent-color (parvocellular) pathway from the retina to the striate cortex rather than "leaking" into the non-opponent (magnocellular) pathway that carries the luminance contribution to motion.[8] These separate routes for color and luminance then converge to form a common motion pathway and a common site for motion aftereffects.[29] Although color does contribute to the second-order attributes as shown in
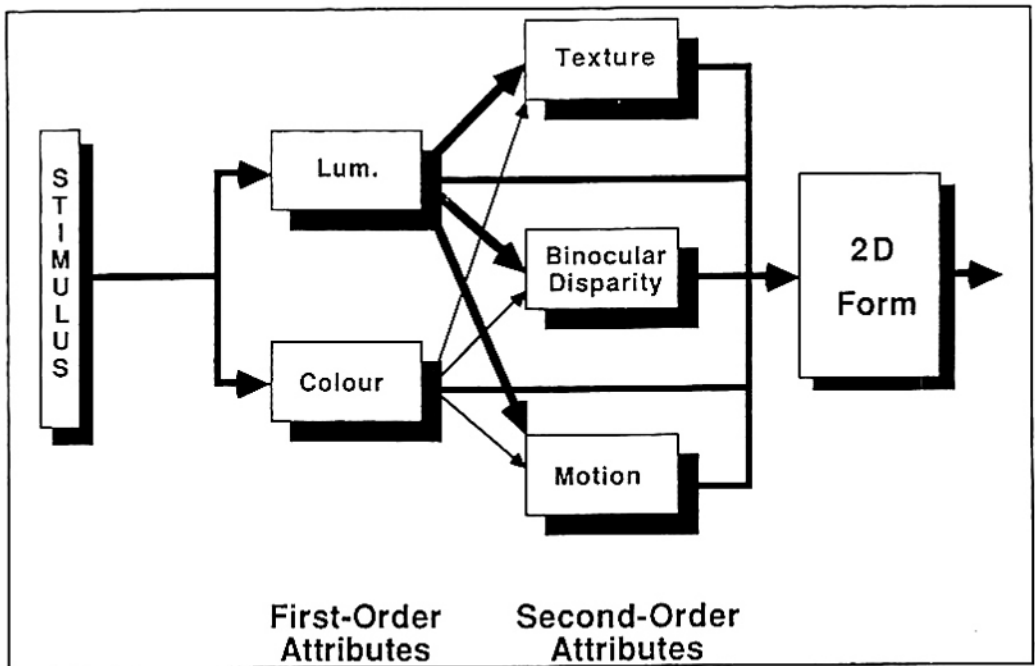
Figure 3. Texture, binocular disparity, and motion are second-order attributes of patterns defined by luminance or color. Color makes a weaker contribution to these second order attributes than does luminance.

Figure 3, its contribution is much weaker than that of luminance.

Shape primitives. The two-dimensional shape of a stimulus is coded on a global level by the spatial arrangement of image information in each pathway. To segment image areas defined by a particular attribute, the visual system must be able to distinguish one region from another which it can do very well by simply coding the value of the attribute (say, color) at each point in a retinotopic map. In the case of luminance, however, it has been demonstrated that the visual system goes beyond this simple coding and extracts local structure (shape primitives) directly with cells that are selective to orientation and size[34] and perhaps curvature.[35] The visual system may use similar shape primitives (receptive-field structures)

for all attributes or may have some specialized encoding for particular attributes, for example, extracting local orientation and size for luminance and but only coding values point by point for other attributes. The shape primitives available for each attribute limit, and in a way, identify the algorithms the visual system can use to represent shapes. For example, position-, size-, and orientation-invariant descriptions that could form the basis for memory and recognition operations[36,37] require size and orientation coding at an early level.

Physiological studies have identified receptive-field structures at several stages in the visual system. In the retina and lateral geniculate, for example, information is coded by antagonistic center and surround organization
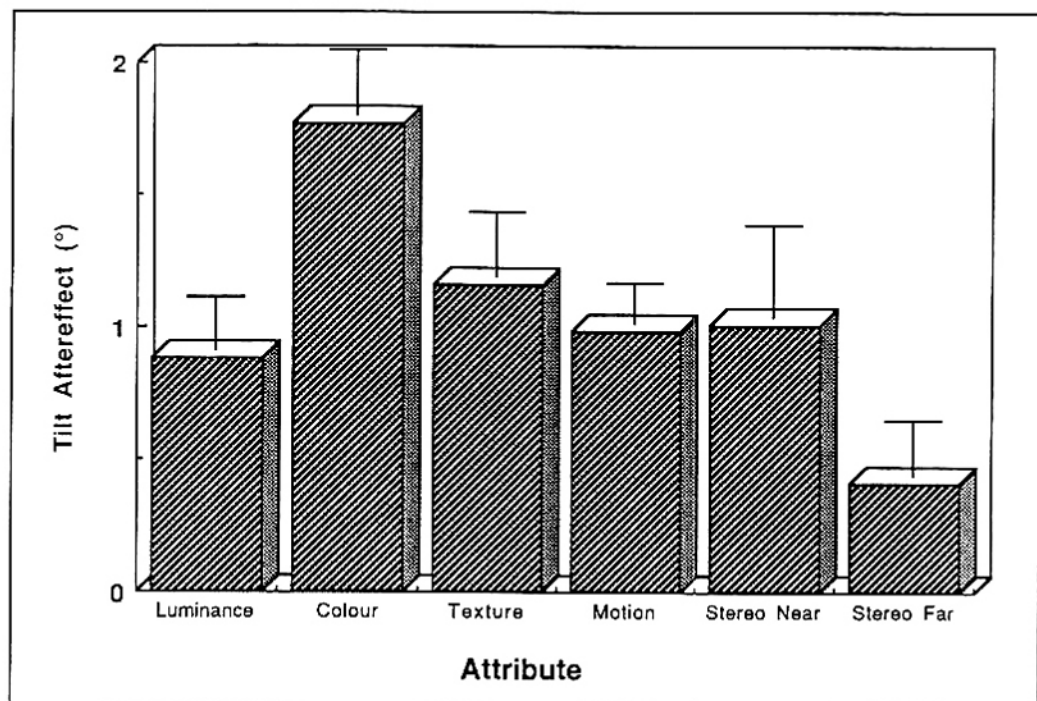
Figure 4. Tilt aftereffects as a function of the attribute defining the bars of the adapting and test gratings. Vertical handles show standard errors (+1.0 S.E.).

and these contribute, in the striate cortex, to the formation of orientation- and size-selective cells.[34] Cells in area V4 appear to have both oriented and non-oriented receptive fields selective for color[3] while many cells in area MT,[38] although directionally selective, do not appear to be orientation-selective, at least not for the orientation of a moving bar. Whether or not they are selective for the orientation of bars defined by relative motion has not been determined. Cells in area V1 and V2 that respond to random-dot stereograms do not appear to be selective for the orientation of bars presented as random-dot stereograms.[39] Although this catalogue of receptive-field structures is extensive, it is not sufficient to identify the coding dimensions ultimately available for all of the attributes we are studying. We

have, therefore, developed a series of psychophysical tests to identify coding primitives for these attributes.

Aftereffects. Size[40] and tilt aftereffects[41] have been used to infer the existence of size and orientation coding dimensions and we have examined these aftereffects for each attribute. We have already demonstrated size aftereffects for color stimuli.[9] Elsner[42] has demonstrated orientation specificity for color stimuli using a tilt aftereffect, and Tyler[43] has reported tilt and spatial frequency aftereffects for random-dot stereograms.

With Patrick Flanagan and Olga Favreau, I have examined the tilt aftereffect paradigm to determine whether there are orientation-tuned detectors specialized for luminance, color, texture, relative motion, and binocular
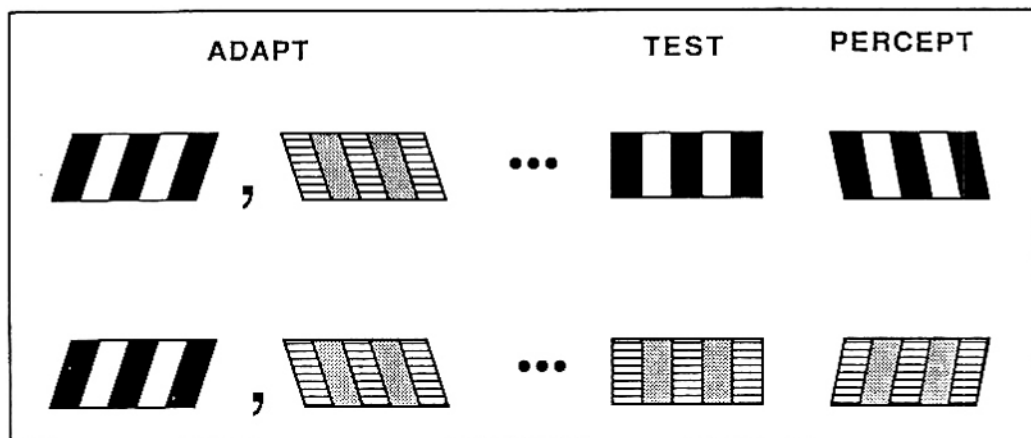
Figure 5. Opposing aftereffects paradigm. Observers alternated adaptation between a luminance-defined grating tilted one way and a texture-defined grating tilted the other way. Vertical tests defined by luminance appeared tilted the opposite direction from the luminance-defined adapting gratings and vertical tests defined by texture appeared tilted the opposite direction from the texture-defined adapting grating.

disparity. We use a standard induction and test procedure for all the candidate pathways: observers are exposed to the adapting stimulus for 8.0 seconds and then to the test for 0.3 seconds in a repeating cycle. The stimuli are square-wave gratings of 0.5 cycles per degree presented in an 8° square display. The adapting stimuli are tilted 15° off vertical. When the test is present, they match the apparent tilt of the test by adjusting a comparison stimulus presented in an unadapted region of the visual field. We recently compared tilt aftereffects for all five attributes (Figure 4) and found that they were all of similar strength.[12,44]

This result may imply either that each attribute undergoes a separate but similar orientation analysis or that a single orientation analysis operates on a higher image that recombines the images from the separate attributes in some manner. In order to distinguish between these two possibilities, we have used an opposing adaptation technique

where we induce opposite aftereffects for two attributes simultaneously. Observers alternately adapt to, for example, texture gratings tilting to the left of vertical and luminance gratings tilting to the right (Figure 5). Three observers participated in these experiments. Following adaptation, vertical test gratings appear tilted to the right if they are defined by texture but tilted to the left if defined by luminance. These independent tilt aftereffects demonstrate that similar analyses of orientation must be occurring in parallel for texture and luminance. In addition, we have been able to show independent aftereffects for color and texture, and color and luminance (Figure 6). In fact, we have found[10] independent analyses of orientation for a luminance pathway and two chromatic pathways, one red-green and the other aligned with the short-wavelength cone axis (tritanopic confusion line). Since orientation analysis does not emerge until the first cortical visual area, this represents evidence of the use
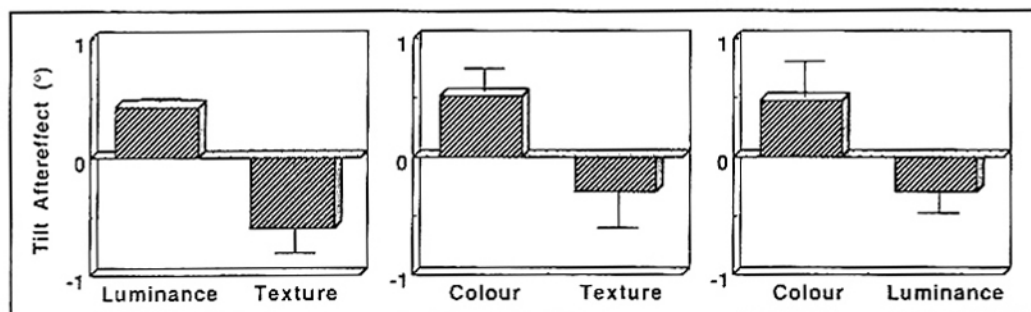
Figure 6. Tilt aftereffects measured in opposing adaptation conditions for luminance versus texture, color versus texture, and color versus luminance. Vertical handles show standard errors (+1.0 S.E.).
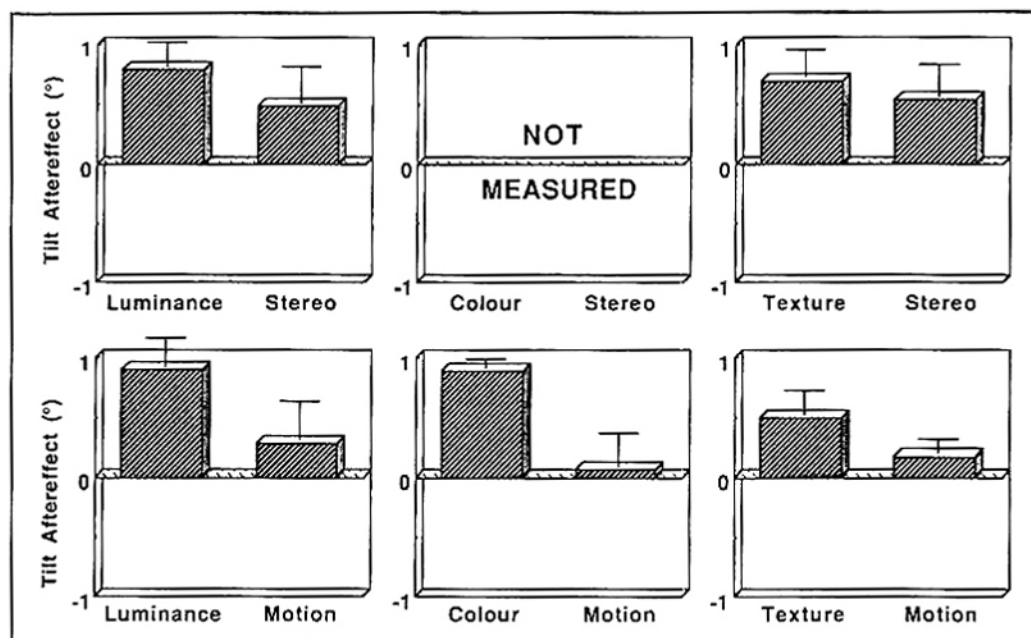

Figure 7. Tilt aftereffects measured in opposing adaptation conditions for stereo versus luminance and texture, and for motion versus luminance, color, and texture. The condition of stereo versus color could not be measured because red-green anaglyphs were used to present the random-dot stereograms and the color stimuli could not be seen properly through the filter glasses. Vertical handles show standard errors (+1.0 S.E.).

of cardinal color axes[45] at the cortical level.

In contrast to these independent tilt aftereffects, we find no independence in the opposed adaptation paradigm whenever one of the tilted gratings is defined by either motion or binocular disparity (Figure 7). In this case, the observed tilt on the motion or binocular disparity-defined test grating is in the same direction as that seen on the other test (luminance-, color-, or texture-defined). This may imply that the adaptation effects of motion or binocular disparity-defined stimuli are simply overwhelmed by the stronger stimuli. If this were the case, then pitting motion-defined stimuli against
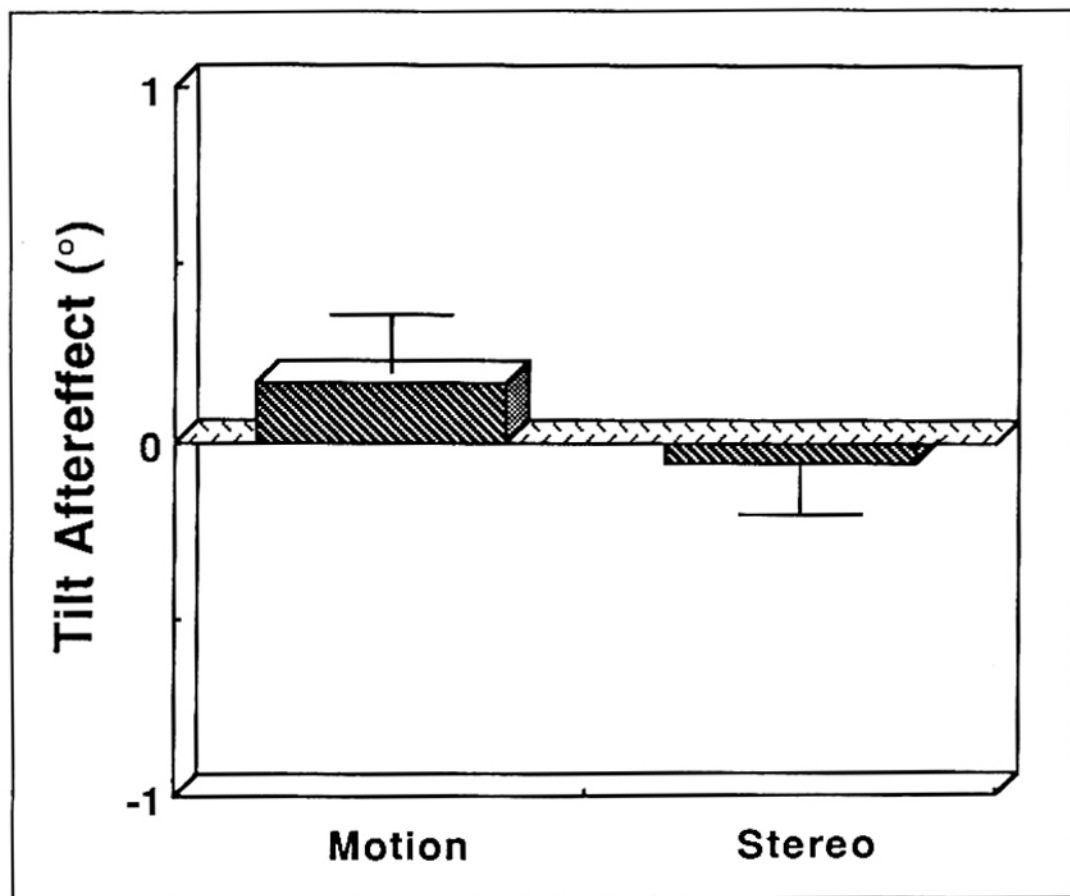
Figure 8. Tilt aftereffects measured in opposing adaptation conditions for motion versus stereo. Vertical handles show standard errors (+1.0 S.E.).

binocular disparity-defined stimuli should involve equal adaptation strengths, and the independent aftereffects should reappear. Figure 8 shows that this is not the case. When motion-defined stimuli and binocular disparity-defined stimuli are opposed, there is little or no aftereffect in either direction. These results suggest that there is no independently adaptable orientation analysis for motion-defined or binocular disparity-defined stimuli. The tilt aftereffects seen in the original experiment following adaptation for motion-defined or binocular disparity-defined stimuli presented alone (Figure 4) must have been based in a higher-level representation, common to all the attributes, that also represented orientation explicitly and was adaptable. For size aftereffects, we have demonstrated independent size aftereffects for color and luminance[9] and we shall shortly begin to test for independence with other attributes.

Illusions. Two illusions, horizontal-vertical illusion and tilt illusion (Zöllner), were tested with composite figures (Figure 9). Figure 10 gives an example of the horizontal-vertical illusion presented as a composite figure with the horizontal bar defined by texture and the vertical bar defined by binocular disparity (the two images in
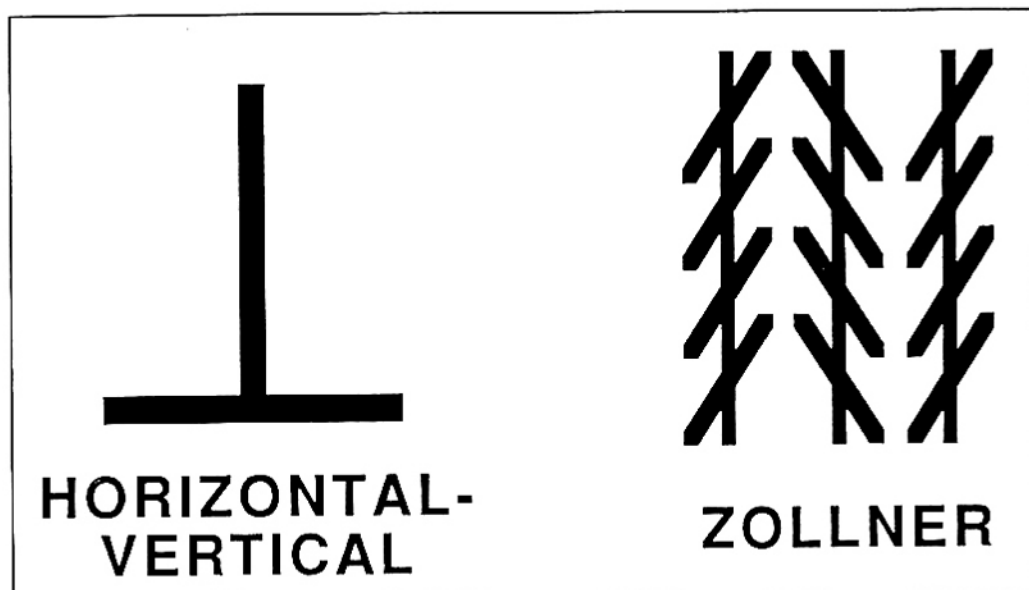
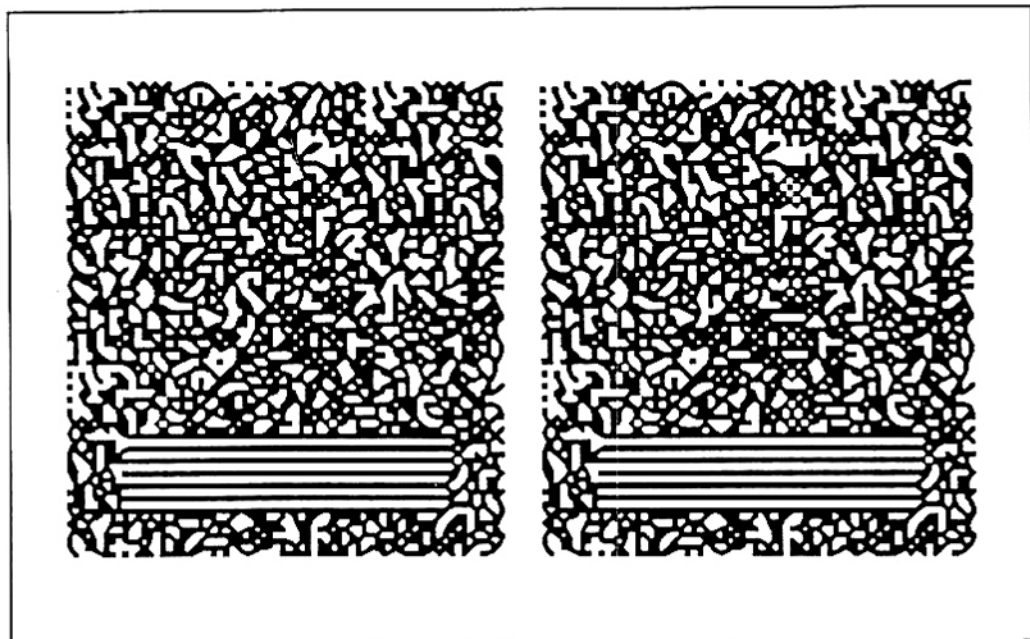**Figure 9.** Two illusions, horizontal-vertical and Zöllner, used to test effects in composite figures.

Figure 10. Horizontal-vertical illusion presented as a composite figure. The two images must be fused to see the vertical bar defined by stereo. The horizontal bar is defined by texture.

the figure must be fused to produce the stereoscopic effect). Composite figures like that of Figure 10 allow us to test the level at which illusions occur. An illusion that remains undiminished in a composite figure must involve pro-cesses that operate on a combined image. An illusion whose strength is reduced in a composite figure must involve processes that operate independently on individual attributes.
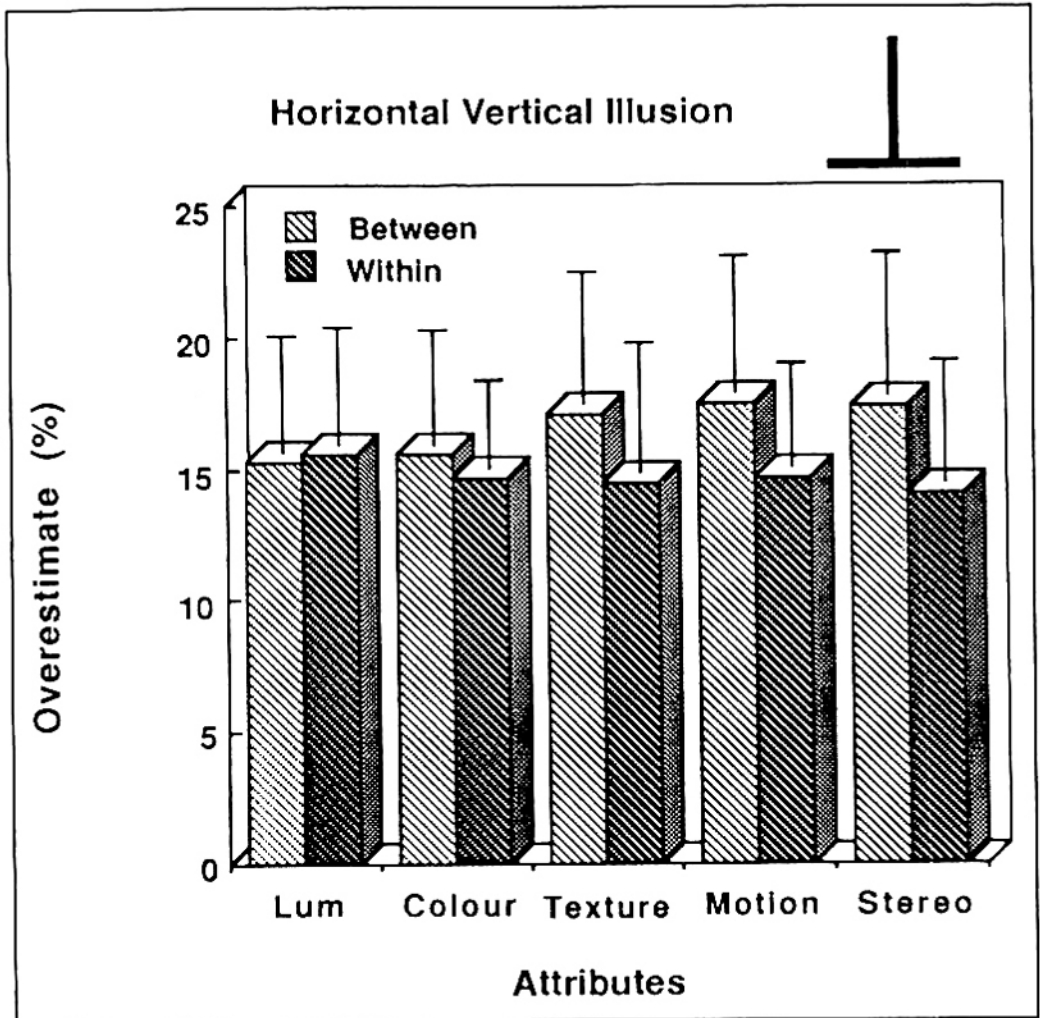
With two exceptions (color and

Figure 11. Horizontal-vertical illusion as a function of stimulus attribute and whether the attribute was used to defined both bars ("within" condition: dark striped columns) or only one of them ("between" condition: light striped columns), the other being defined by a different attribute. The vertical axis shows the percent lengthening of the horizontal bar required for it to appear equal in length to the vertical bar. Vertical handles show standard errors (+1.0 S.E.).

stereo could not be combined because red-green anaglyphs were used to present the random dot stereograms, and the color stimuli could not be seen properly through the filter glasses), all possible combinations of the five attributes were used in these figures producing 23 different versions: five "within" figures with both the inducing and the test portion of the illusion figure defined by the same attribute and 18 "between" figures with the two components defined by different attributes. The display subtended 10° of visual angle and, for the horizontal-vertical illusion, observers adjusted the length of the horizontal bar until it appeared to equal that of the vertical bar. Four observers participated. The results are shown in Figure 11. The values shown for the "between" results for each attribute are the average of results involv-
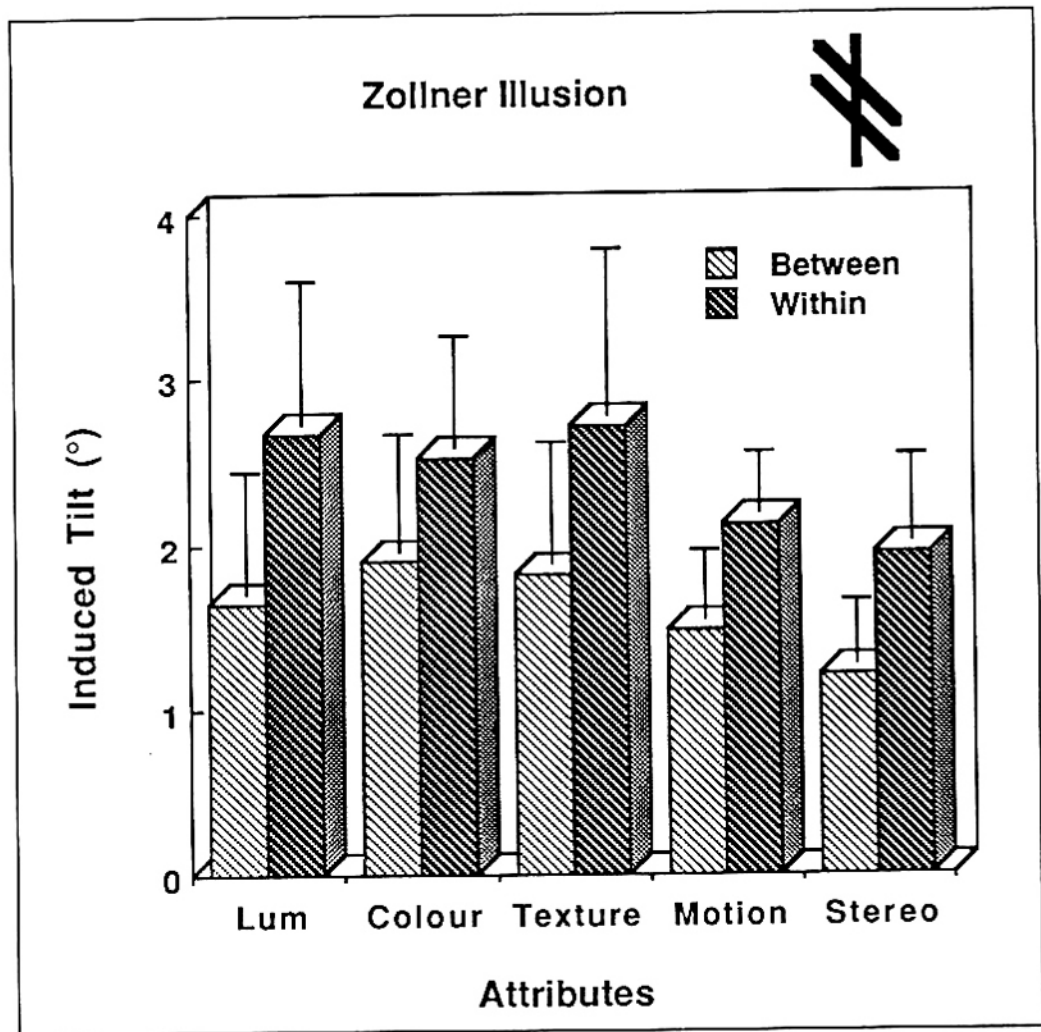
Figure 12. Zöllner illusion as a function of stimulus attribute and whether the attribute was used to define both the central shaft and the inducing bars ("within" condition: dark striped columns) or only one of them ("between" condition: light striped columns), while the other was defined by a different attribute. The vertical axis shows the apparent angular tilt of the central shaft. Vertical handles show standard errors (+1.0 S.E.).

ing that particular attribute in all its combinations with the four others. The strength of the horizontal-vertical illusion was very similar for all the attributes whether both lines were defined by the same attribute (within) or by different attributes (between).

For the Zöllner illusion, a small dot was placed 2° from the end of the central shaft of the Zöllner figure. Observers adjusted the position of this dot until they felt that it was colinear with the apparent direction of the central shaft. Four observers participated. Other than the change in the illusion figure, the conditions used were the same as those in the horizontal-vertical tests described above. The results are shown in Figure 12. The strength of the illusion was similar for all attributes in the
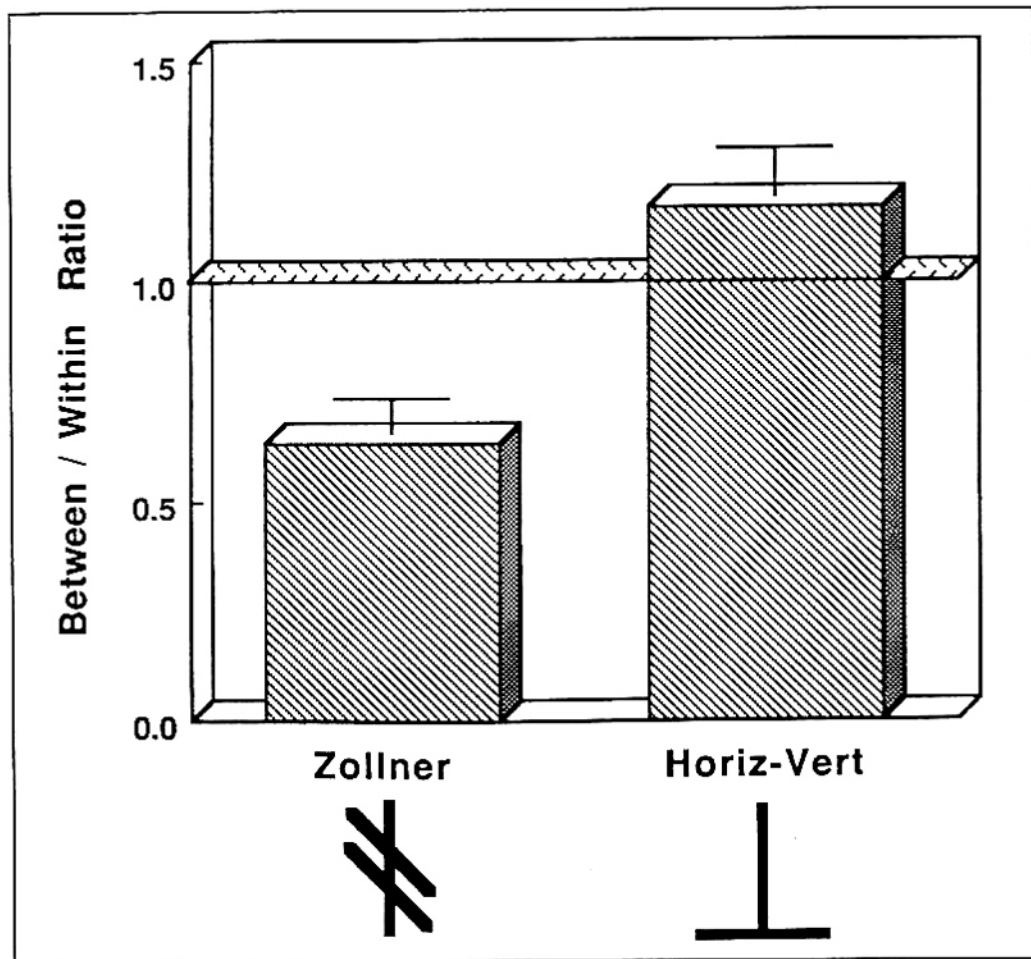
Figure 13. Illusion strength for "between" conditions divided by that for "within" conditions as a function of illusion. Vertical handles show standard errors (+1.0 S.E.).

"within" conditions and similar, as well, across attributes in the "between" conditions. However, the illusion strength was uniformly smaller in the "between" conditions than in the "within" conditions.

A "between-within" ratio was computed by dividing the average illusion strength for all "between" figures by the average strength for all "within" figures (Figure 13). A value of 1.0 for this ratio indicates that processes responsible for the illusion could access a combined shape representation while a value less than 1.0 indicates that the processes responsible must, at least in part, be located in the individual pathways so that the integrated figure was unavailable to them and so could not trigger the illusion. The horizontal-vertical illusion had a between-within ratio that did not differ significantly from 1.0, indicating that the processes underlying the illusions access an integrated image. Processes involved in spatial scaling of horizontal and vertical dimensions might be expected to operate on an integrated image since

Figure 14. Schematic representation of the difference in tilt illusion for a figure with shaft and inducing bars both defined by the same attribute (top) and the same figure with the shaft defined by one attribute and the inducing bars defined by another (bottom). On the top in the "within" condition, a tilt illusion is generated by orientation analyses both in the independent pathways depicted in the center and the common representation at the right. The two effects combine to produce a total effect greater than either alone. On the bottom in the "between" condition, a tilt illusion is generated by orientation analyses only in the common representation at the right. The inducing bars and the central shaft are represented in separate independent pathways and so no interaction can take place to produce a tilt illusion at this level. The final result in the "within" condition is therefore larger than in the "between" condition.
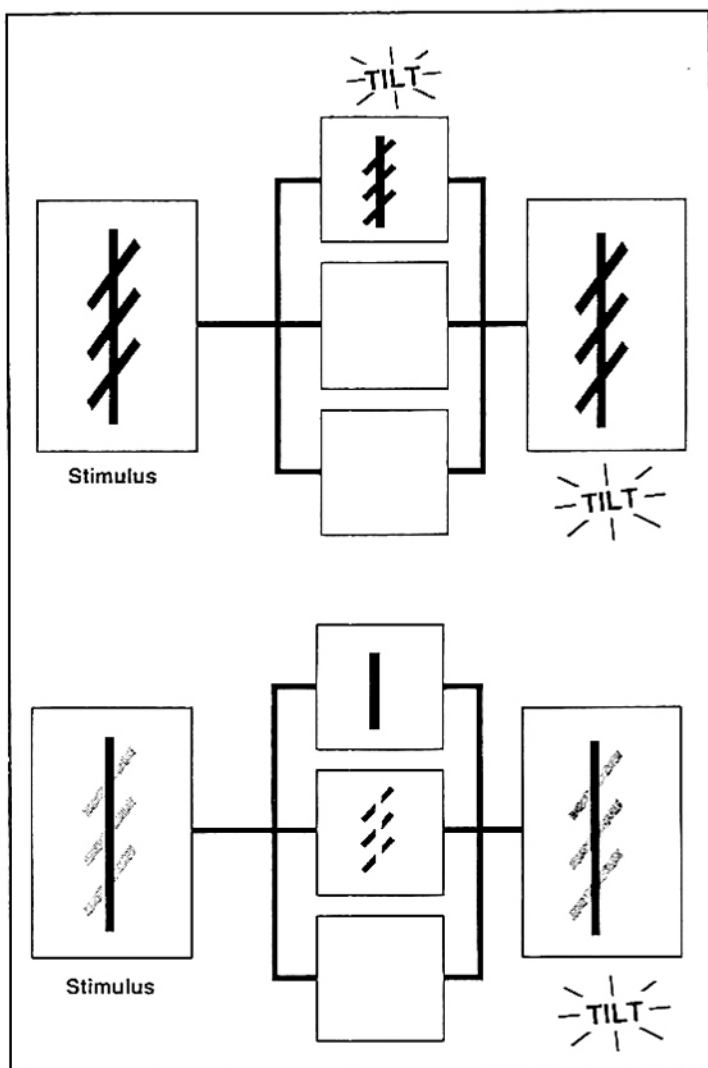
they must deal with scenes made up of objects defined by an arbitrary assortment of attributes. For the Zöllner illusion, on the other hand, the ratio was significantly less than 1.0, implying that the orientation coding underlying this tilt illusion must be occurring independently in each pathway to some extent: a composite figure would, therefore, produce less illusion than one defined by a single attribute (Figure 14). The ratio was, nevertheless, also greater than zero indicating that there is some interpathway interaction for orientation coding or alternately some orientation coding in a common high-level representation.

The previous experiments on tilt aftereffects had indicated that stimuli defined by motion or binocular disparity did not produce independent tilt aftereffects. For purposes of comparison, the results of the tilt illusion for figures involving motion and binocular dispar-
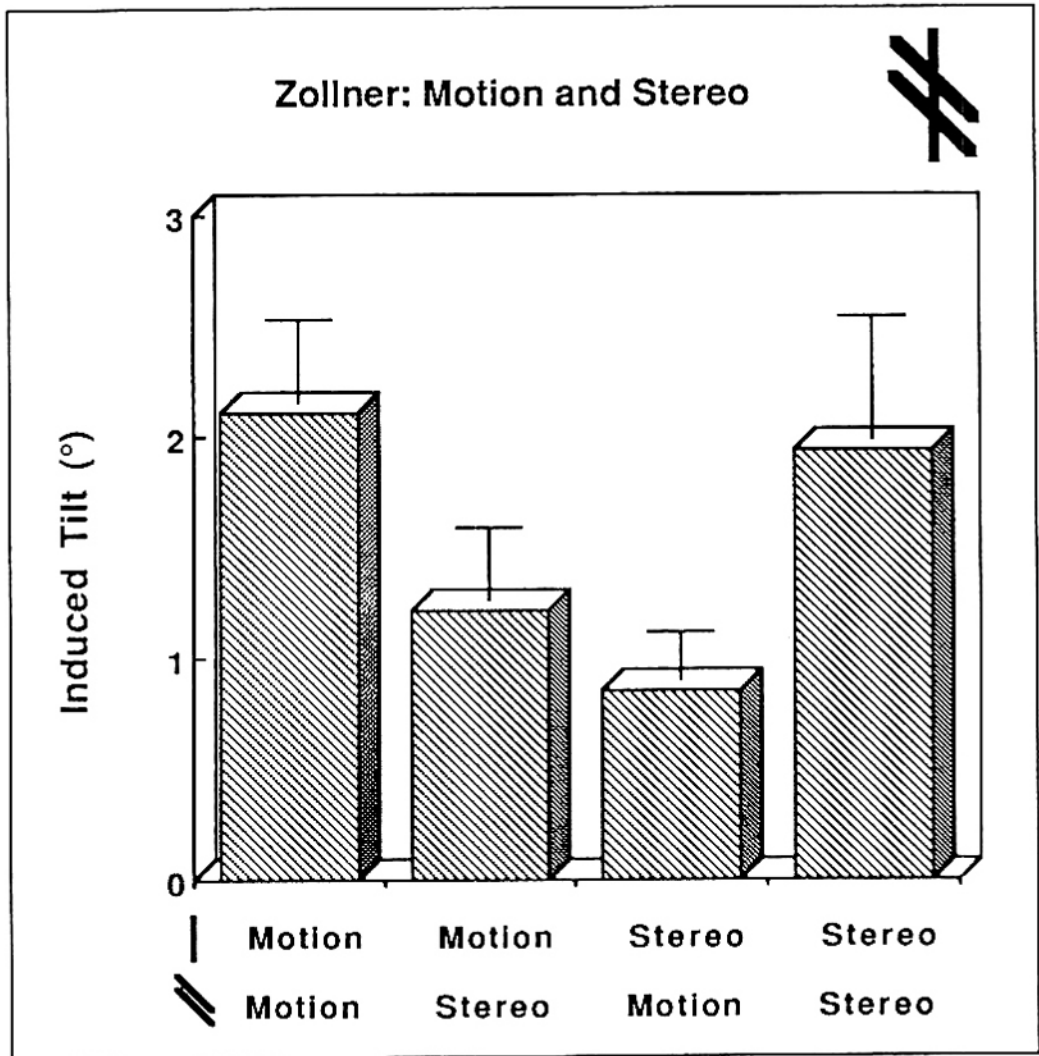
Zollner: Motion and Stereo

Figure 15. Zöllner illusion as a function of the stimulus attribute defining the central shaft and that defining the inducing bars for the attributes stereo and motion only. The vertical axis shows the apparent angular tilt of the central shaft. Motion-motion and stereo-stereo are "within" conditions while motion-stereo and stereo-motion are "between" conditions. Vertical handles show standard errors (+1.0 S.E.).

ity are shown in Figure 15. The data show a drop in illusion strength for Zöllner figures composed of both motion and binocular disparity-defined bars compared to those composed of only motion-defined bars or only binocular disparity-defined bars. The tilt illusion data suggests that there is independent analysis of orientation for motion binocular disparity pathways, at least to the same extent as for the other pathways. The results from the tilt aftereffect study indicate that the orientation analysis, although it occurs, may not involve adaptable mechanisms. Both studies do suggest the role of an orientation analysis that operates on a common representation.

Discussion

These results raise several intriguing questions: why there are several analyses, why they seem similar and why they involve, in particular, size and orientation. Since our work here is not yet complete, the answers are speculative and offer interesting directions for further research.

Why are there several analyses? It is clear that if at least one analysis of shape, say based on luminance, is useful, more may be significantly better. Many surfaces are distinguished from their surrounds not only by luminance differences but also by texture and color so that shape analyses for additional features would improve the segmentation of the surfaces in the scene. In fact, it is easy to argue that luminance, the favored attribute in most shape analysis programs in computer vision, is a problematic choice because many luminance edges in the scene are irrelevant shadow borders. Color and texture differences are much more reliably linked to object boundaries than are luminance edges.

Why are they similar? We have established that there is independent orientation analysis for several attributes and independent size analysis for at least two. We have not yet established that the analysis of each attribute is similar in every respect but one very important reason for the degree of similarity that we have found may be the exchange of image information between analyses. If an integrated higher-level image is to be formed, it is useful to have all lower-level images defined in a standard format. Size and orientation coding may therefore be part of an internal standard for image exchange in the visual system.

If the shape codes for all the attributes are similar, we would expect that processes that depend on two-dimensional shape should be effective on a given two-dimensional shape no matter what attribute is used to define it. For example, we should be able to interpret an ellipse as a tilted circle whether the ellipse is defined by color (red on green of equal luminance), by a random-dot stereogram or by luminance (black on white). Phenomena that depend on information other than shape or that depend on a particular encoding of two-dimensional shape may be preserved for stimuli defined by some attributes but lost for others.

In 1971, Julesz[46] reported that shapes defined by random-dot stereograms produced classical visual illusions, identifiable letters and various other perceptual phenomena. In extending Julesz's approach to additional attributes, we have found that all three-dimensional shape inferences involving objects defined by explicit contours work for all attributes.[18,19,21,22,47] These results are directly opposed to the claims of Livingstone and Hubel[14] that monocular depth cues are ineffective unless luminance defines the stimulus contours. Our studies clearly show that this is not the case, and we attribute their failures to find three-dimensional depth to the fine detail in their images, detail that is critical to the interpretation of their figures. In their stimuli, this fine detail can only be resolved across the entire image if luminance is present

and both the stimulus and the depth are lost at equiluminance.

Not all three-dimensional inferences are possible for every attribute, however. Stimuli involving implicit object borders appear to require luminance. For example, shadows and subjective contours both can be perceived with stimuli defined by luminance but both fail to be perceived for the same stimuli when defined by other attributes (for example, Figure 3) even when the object is in motion.[47] The luminance pathway is essential for shadows and subjective contours and the essential information that is derived in the luminance pathway may be the polarity of contrast across the borders.[19,48]

Why size and orientation? There are two possible answers for the popularity of these two image codes. The first is that local size and orientation information significantly improves the detection of contours in the image. These multi-scale approaches[49] to contour identification are the initial stages of contour labeling involved in recognition-by-component models. Contours are extracted and object parts, whether generalized cylinders[23] or geons,[24] are identified from these contours using simple image constraints. These early representations then index memory to retrieve object identity and other useful information for completing the three-dimensional model. Memory representations are assumed to be object-centered in that arbitrary views of the object can be matched to the memory prototype.

The second possibility is that size and orientation codes are an intermedi-ate step in obtaining size- and orientation-invariant shape descriptions. Fourier-Mellin transforms[50] and Fourier-log polar-Fourier transforms[36,37] are representatives of this class of transform. The coincidence between the size and orientation representation which is an intermediate step in these transforms and that seen in local regions of striate cortex, and revealed by the aftereffects that we have studied, may simply be a coincidence. But it may be an indication that the visual system uses a transform of this type.

Cortical areas. It has been proposed that the existence of separate cortical areas specialized for different visual attributes allows for a more economical analysis of each. Although each of these areas may perform some specialized tasks such as color constancy or contour polarity, the results of the experiments reported here suggest that, other than these specialized analyses, the independent analysis of visual attributes may involve only the extraction of two-dimensional shape. The cortical areas that have been proposed as the sites for these independent analyses include all those up to MT and V4.[3] These independent representations would involve a very large proportion of visual cortex in what is only a preliminary step in vision. It may be that the functionally independent analyses that we have demonstrated in our experiments occur before these regions and involve perhaps only compartmentalized representations rather than physically separate ones. Orientation tuning for disparity- and direction-selective units is seen as early as the stri-

ate cortex. However, in our oriented gratings presented as random-dot stereograms or random-dot kinematograms, these striate units must respond to the orientation of individual, luminance-defined elements and not to the orientation of the alternating broad strips of random dots which are at different depths (stereogram) or moving in different directions (kinematogram) and which all have the same mean luminance. Some of the orientation effects we measure must therefore be based in extrastriate cortices although not necessarily as far along as V4 or MT. Area MT and V4 might then be candidates for a new set of pathways that are not specific to stimulus attributes. V4 might be the site of a common two-dimensional shape representation that combines information from all visual attributes. Note that if this were the case, V4 cells would not necessarily respond to motion, for example, but to the shapes defined by motion differences between regions; it would not necessarily respond to stereo but to the shapes of regions defined by stereo. Specifically, some cells in V4 ought to show tuning for the orientation of contours that are defined by any feature—color, luminance, texture, motion, or disparity—as if these cells performed an OR function on the outputs of lower-level, orientation-tuned units. Conversely, area MT might respond to motion of shapes no matter how they are defined.

Acknowledgments

References

1. A. Cowey, *Q. J. Exp. Psych.* **31**, 1 (1979).
2. H.B. Barlow, *Vision Res.* **26**, 81 (1986).
3. S.M. Zeki, *Nature (London)*, **274**, 423 (1978).
4. D.C. van Essen, in *Cerebral Cortex Vol. III*, A. Peters and E.G. Jones, Eds., (Plenum Publishing, New York, 1985), p. 259.
5. J. Zihl, *Brain* **106**, 313 (1983).
6. A. Damasio *et al.*, *Neurology* **30**, 1064 (1980).
7. J. Rovamo, L. Hyvarinen, and R. Hari, *Doc. Ophthal. Proc. Series* **33**, 457 (1982).
8. J.H.R. Maunsell and W.T. Newsome, *Ann. Rev. Neurosci.* **10**, 363 (1987).
9. O.E. Favreau and P. Cavanagh, *Science* **212**, 831 (1981).
10. P. Flanagan, P. Cavanagh, and O.E. Favreau, *Optics News*, **13**, 161 (1987).
11. P. Flanagan, P. Cavanagh, and O.E. Favreau, *Invest. Ophthalmol. Vis. Sci. Suppl.* **29**, 327 (1988).
12. P. Cavanagh, paper presented to the Symposium on New Insights on Visual Cortex, Rochester, June 1988.
13. M.S. Livingstone and D.H. Hubel, *J. Neurosci.* **4**, 309 (1984).
14. M.S. Livingstone and D.H. Hubel, *J. Neurosci.* **7**, 3416 (1987).
15. H.H. Bülthoff and H.A. Mallot, *Invest. Ophthalmol. Vis. Sci. Suppl.* **29**, 40 (1988).
16. G. Sperling and B. Dosher, *Vision Res.*, **26**, 973 (1986).
17. R.L. Gregory, *Proc. R. Soc. London*, B**204**, 467 (1979).
18. P. Cavanagh, *J. Opt. Soc. Am.* A**2**, P51 (1985b).
19. P. Cavanagh, *Bull. Psychonomics Soc.* **23**, 273 (1985c).
20. P. Cavanagh, *Bull. Psychonomics Soc.* **24**, 291 (1986).
21. P. Cavanagh, *Computer Vision, Graphics and Image Processing*, **37**, 171 (1987).
22. P. Cavanagh, in *Computational Pocesses in Human Vision: An Interdisciplinary Perspective*, Zenon Pylyshyn, Ed. (Ablex, Norwood, N.J., 1987), p. 254.
23. D. Marr, *Vision* (Freeman, San Francisco, 1982).
24. I. Biederman, *Psych.Rev.* **84**, 115 (1987).
25. P. Cavanagh, S. Peters, and M. von Grünau,

paper presented to the annual meeting of the European Conference on Visual Perception, Bristol, August 1988.

26. P. Cavanagh and M. von Grünau, paper presented ARVO, Sarasota, May 1989.

27. P. Cavanagh, J. Boeglin, and O.E. Favreau, *Perception* 14, 151 (1989).

28. P. Cavanagh, C.W. Tyler, and O.E. Favreau, *J. Opt. Soc. Am.* 1, 893 (1984).

29. P. Cavanagh and O.E. Favreau, *Vision Res.* 25, 1595 (1985).

30. P. Cavanagh and S.M. Anstis, *Invest. Ophthalmol. Vis. Sci. Suppl.* 27, 291 (1986).

31. C.M.M. de Weert and K.J. Sadza, in *Colour vision: Psychophysics and physiology*, J. D. Mollon and L. T. Sharpe, Eds. (Academic Press, London, 1983), pp. 553-562.

32. D.L. Grinberg and D.R. Williams, *Vision Res.* 25, 531 (1985).

33. P. Cavanagh, *Invest. Ophthalmol. Visual Sci. Suppl.* 29, 196 (1988).

34. D.H. Hubel and T.N. Wiesel, *J. Physiol.*, 195, 215 (1968).

35. A. Dobbins, S.W. Zucker, and M.S. Cynader, *Nature (London)* 322, 196 (1987).

36. P. Cavanagh, in *Figural Synthesis*, P. C. Dodwell and T. Caelli, Eds. (Lawrence Erlbaum Associates, Hillsdale, N.J., 1984), pp. 185-218.

37. P. Cavanagh, in *Models of the Visual Cortex*, D. Rose and V.G. Dobson, Eds. (John Wiley & Sons, London, 1985), pp. 85-95.

38. J.A. Movshon *et al.*, in *Pattern Recognition Mechanisms*, C. Chagas, R. Gattass, and C. Gross, Eds. (Springer Verlag, Berlin, 1986), pp. 117-151.

39. G.F. Poggio *et al.*, *Vision Res.* 25, 397 (1985).

40. C. Blakemore and P. Sutton, *Science* 166, 245 (1969).

41. F.W. Campbell and L. Maffei, *Vision Res.* 11, 833 (1971).

42. A. Elsner, *Perception and Psychophysics* 24, 451 (1978).

43. C.W. Tyler, *Perception* 4, 187 (1975).

44. P. Cavanagh, M. Arguin, and P. Flanagan, paper presented to the annual meeting of the English Pychological Association, Oxford, July 1987.

45. J. Krauskopf, D.R. Williams, and D.W. Healy, *Vision Res.* 22, 1123 (1982).

46 B. Julesz, *Foundations of Cyclopean Perception* (University of Chicago Press, Chicago, 1971).

47. P. Cavanagh and V.S. Ramachandran, paper presented to the annual meeting of the Canadian Psychology Association, Montreal, June 1988.

48. P. Cavanagh, and Y. Leclerc, *J. Exp. Psychol.[Hum. Percept.]*, 3 (1988).

49. A.P. Witkin, *Proceedings of the 8th International Joint Conference on Artificial Intelligence* 2, 1019 (1983).

50. D. Casasent and D. Psaltis, *Applied Optics* 15, 1793 (1976).