Contour Completion at Edge Endings

by

Joseph Edward Scheuhammer

B.A.(Hon) University of Western Ontario (1978)
Philosophy and Psychology

M.A. University of Western Ontario (1979)
Psychology

Submitted to the Department of
Brain and Cognitive Sciences
in Partial Fulfillment of the Requirements
for the Degree of

DOCTOR OF PHILOSOPHY
in
COGNITIVE SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June, 1988

Signature of Author ....
Department of Brain and Cognitive Sciences
February 5, 1988

Certified by ....
Professor Jeremy Wolfe
Department of Brain and Cognitive Sciences

Accepted by ....
Professor Emilio Bizzi
Head, Department of Brain and Cognitive Sciences

Contour Completion at Edge Endings

Joseph Edward Scheuhammer

Submitted to the Department of Brain and Cognitive Sciences
on February 5, 1988
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

Abstract

The contours of a visual scene are not always projected in their entirety into the image. The result is a collection of edges separated by gaps. This thesis is concerned with the manner in which edges end and how that affects contour completion. Edges end in two ways: suddenly, or by gradually fading away. The latter situation implies continuity and provides a clue that the underlying scene contour does not end. When the edges end suddenly, however, the discontinuity implies that the corresponding scene boundary does end, and the gap should not be filled in. Furthermore, even if the edges end smoothly, if there are other discontinuities nearby, these should interrupt completion.

These hypotheses were tested in two ways. Using a gap detection task, subjects were presented with gapped and complete bars, flashed briefly, and were required to judge whether the stimulus was gapped or not. Subjects tended to judge a stimulus as "gapped" if the gap was bounded by abrupt as opposed to smooth endings. Also, when luminance discontinuities cued the presence of the gap, the proportion of "gapped" judgments increased, even in the case of complete stimuli.

Contour completion was further investigated using the Poggendorff Illusion. It was found that gaps bounded by smooth endings induced a significantly greater effect compared to those bounded by abrupt endings. In addition, when the endings were abrupt, the effect was modulated by the angular tilt of the top and bottom edges of the gap. This did not occur when the endings were smooth, suggesting that the gap was perceived differently depending on the type of ending. Thus, smoothly ending edges induced Poggendorff effects like those induced by luminance edges; abrupt endings did not.

The problem of contour perception based on the intermittent edges found

in images is solved, in part, by noting the manner in which the edge ends. If it ends smoothly, and there is no other information to cue that ending, then it is a candidate for extension. If it ends suddenly, then that ending ought to be taken as given.

Thesis Supervisor: Professor Jeremy Wolfe
Professor, Department of Brain and Cognitive Sciences

Acknowledgments

During the seeming age it took to do this thesis, I have had the good fortune to meet and know many people who have helped, in their way, to make it come about. I would like to take this opportunity to express my gratitude to them.

Although they had nothing to do with the research, I must acknowledge Ian, Elizabeth, Howard, Dean, Judy, John, Scott and Tony for many imaginative adventures. Without the diversion, I *might* have finished sooner, but life would have been far more dull.

I thank Joe Bauer, specifically for relating how a certain statistician had scaled the preference of nine vegetables using the method of paired comparisons, and, generally, for conversations on many topics, both relevant and irrelevant.

Thanks to Brian Timney and Zenon Pylyshyn for being there with an umbrella when it started to rain.

A special note of gratitude to Magi (David Wiseman), a Unix wizard who, above and beyond the call of duty, was always willing to patiently explain some obscure facet of C and/or the Unix kernel. Without his help, it would have taken far longer to coerce the apparatus into generating the displays and running the experiments. "Do go unto wizards for advice, for they are cunning, and slow to anger."

I thank the members of my committee—Professors Patrick Cavanagh, Dick Held, Molly Potter, and Jeremy Wolfe—for their thoughtful criticisms and support over the last few months. I am especially indebted to Jeremy who conscientiously read umpteen[1] versions of the thesis. His voluminous red marks, while annoying, were of considerable benefit. He helped me to say what I wanted to in much clearer prose.

I thank my parents, Magdalena and Josef, for their love and support, especially over the last three years.

And, finally, I would like to thank Judy for saying one magic word at those times when it was needed most. Often, I would feel disheartened about the whole endeavour and mumble something along the lines of "... *if* I ever get my degree." She would always counter with these encouraging words: "No, no, no. You mean '*When* you get your degree ... '"

[1]Several, a good many.

Biographical Note

One episode typifies how I will always remember my grandmother. One day, while I was mowing her lawn and she was working in her garden, her next door neighbour leaned over the fence to comment on something—I couldn't hear what it was over the din of the lawnmower. Later, Grandma told me that the topic of conversation had been me, specifically my long hair. Now, this was the mid-sixties when it was still considered radical for males to wear their hair below the tops of their ears. The neighbour apparently thought that there was no hope for society, especially when (formerly) good boys like myself were turning their backs on the old values and progressing down the road to moral ruin. Grandma replied to this indictment by telling the neighour that hair length had nothing whatsoever to do with what was inside a person, and that if I wanted to have long hair that was all right (as long as I kept it clean). The neighbour was non-plussed. Here was an elderly woman, who was presumably part of the "establishment", defending a then radical idea. What I will always remember about this incident is how *rational* Grandma's attitude was. She did not decide what was right or wrong based simply on the accepted mores of the time; she used her mind to reflect on the issues. Other incidents reinforced this impression of Grandma, and I will continue to remember her as something of a critical thinker. I often wonder what she could have accomplished had there been the opportunity to pursue a "higher education".

My own pursuit began in the fall of 1974 as a double major in Philosophy and Psychology at the University of Western Ontario (I used to tell my friends that I studied both normative and applied mental processes). Two people influenced me most during that period. One was William Demopoulas, who taught courses on Epistemology and Philosophy of Logic. In Epistemology, I learned how the logical positivists had erred, but came away with an appreciation of their analytic style. In Philosophy of Logic, I learned of formal systems and, in particular, Turing's seminal work on compuation. I think the most insightful characterization of symbol processing remains his *Computable Numbers.*

That same year, I took a course entitled "Minds and Machines" from one Zenon Pylyshyn. It was my first encounter with Artificial Intelligence and the idea that mental activity is fruitfully viewed as information processing; a view that I think is essentially correct. Under Zenon's tutelage, I eventually became interested in the relationship between perception and visual imagery. One day, towards the end of my undergraduate career, Zenon recommended I apply to graduate school at MIT because, in his words, "you think like they do."

Before I got to MIT, I spent a year working on a Master's thesis with Brian Timney. It was my first exposure to honest-to-goodness perceptual psychology. Brian introduced me to the idea that there was more to the universe than cognition.

When I did arrive here, at the then Department of Psychology, it was with the idea that I was going to do the ultimate imagery experiment. I never got around to it because I came to feel that there are too many loop holes on either side of the issue: A truly crucial experiment was not in the offing. In addition, the claim that imagery and perception were functionally equivalent made me increasingly interested in what perception was all about, so much so that the imagery side of my interests was eventually eclipsed. I took an introductory course in perception from Professors Held and Wolfe. Later, at Molly Potter's urging, I took a course entitled "Natural Computation and Control" taught by Whitman Richards. There I learned of David Marr's approach to vision. Two lessons were gained from this experience: (1) that something as simple as seeing can be extraordinarily complicated when it comes down to explaining it, and (2) that one should spend time on the question of the purpose of the visual system's computations, the "what" and "why" of visual processes. If the visual system feels fit to extract something from the retinal mosaic, one should ask why that should be. What good does it do to process that information? I have tried to pay some attention to this question in the research that follows.

Due to unfortunate circumstances, I had to leave MIT and return to Western between the years 1984–1987. My exile was not without its virtues, however. For one, in the course of putting food on the table, I was enlisted to teach undergraduate courses in the History of Psychology, and Cognitive Science. I discovered that teaching is a very enjoyable and rewarding activity (marking, however, is not). Also, with respect to the history course, I taught from the original works of Wundt, James, Freud, Skinner, and Köhler, whose ideas are much better expressed by themselves than a secondary source. Particularly in the case of Köhler, I was frequently surprised at how insightful the Gestaltists were—textbook accounts do not do the Gestalt movement justice.

Psychological research is, one might say, 90% apparatus. It certainly was in my case. In the course of setting up the lab equipment to run the experiments, I became intimately acquainted with the Unix operating system and discovered the (distracting) joy of computer programming. I hope to make programming a facet of my future.

The most pleasant consequence of my return to Canada was that I met Judy Benger, whom I shall marry in June.

And, oh yes, I managed to do a couple of experiments ...

DEDICATED TO THE MEMORY OF
ELISABETHA SOHL

This research was supported in part by a National Sciences and Engineering Research Council of Canada (NSERC) postgraduate scholarship awarded to Joseph Scheuhammer, NSERC operating grant A7062 awarded to Brian Timney, NSERC operating grant A2600 awarded to Zenon Pylyshyn, and National Institute of Health core grant 5P30-EY02621.

Contents

List of Figures

List of Tables

Chapter 1

Edges and Endings

> ... the idea that extracting edges and lines from images might be at all difficult simply did not occur to those who had not tried to do it. It turned out to be an elusive problem: Edges that are of critical importance from a three-dimensional point of view often cannot be found at all by looking at intensity changes in the image. Any kind of textured image gives a multitude of noisy edge segments and even if an edge has a clear existence at one point, it is as likely as not to fade out quite soon, appearing only in patches along its length in the image (Marr, 1982, p. 16).

When we view the world around us, one of the objects of our perception are the contours in the scene presented to our eyes. For example, we see the boundary between that portion of a surface within shadow and that without; see where one object ends and another begins; and can visually trace the whorls of a wood grain. It is generally believed that contour perception is accomplished by the detection of luminance changes in the retinal image. That is, contours are projected into the image and registered as changes in intensity, or "edges", and the localization of these edges will result in the recovery of scene contours.

But edge detection cannot be the whole story. Frequently the contours in a scene are not projected veridically into the image and do not give rise to luminance changes. Consider Figure 1.1. This is an image of an egg-shaped object in front of a vertical planar surface.[1] Since it is curved, the surface of

[1]The shape is in fact an elliptical paraboloid viewed head on. The equation of this

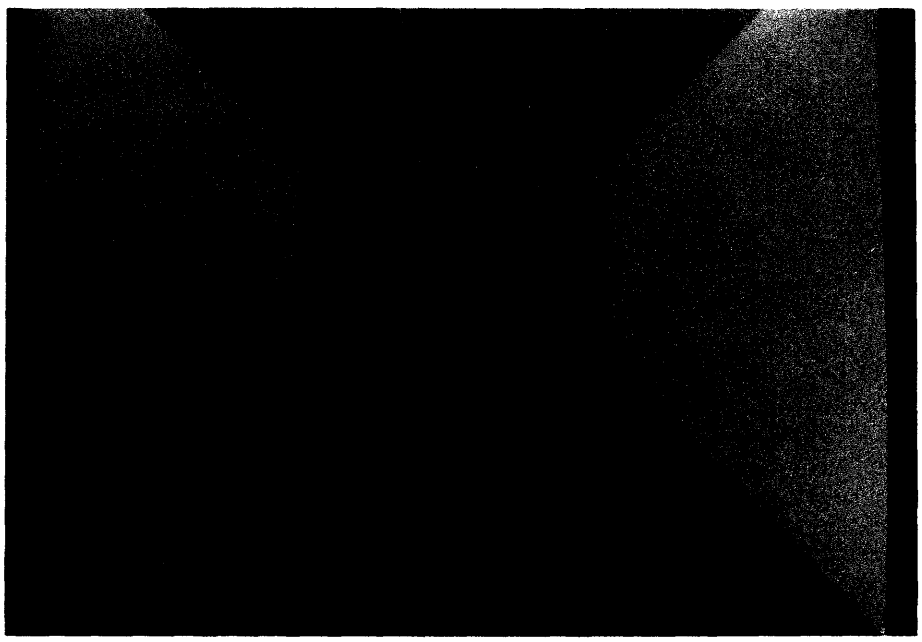

Figure 1.1: A zero contrast contour.

the egg reflects a varying amount of light proportional to the angle between the surface normal and the illuminant direction. Consequently, the surface is depicted in the image as a smooth luminance gradient within the bounds of the occluding contour. The plane behind the egg is of constant luminance, a value specifically chosen such that there is no contrast between part of the egg's surface and its background. Thus, although there is a closed occluding contour in the *scene*, the edge corresponding to it is broken and there is in fact a region where it is non-existent. The extent of the zero contrast region is illustrated in Figure 1.2.

If the image in Figure 1.1 is examined closely, then one observes that, indeed, the occluding boundary vanishes at the top for a short length of the contour. However, if looked at casually, there is a sense that the elliptical edge, corresponding to the occluding contour, does not possess a "gap" and is complete. Somehow, the visual system perceives an edge in a region where there is none and (correctly) joins together the pair of points between which the edge disappears. How this is accomplished is the topic of this thesis: Specifically, what conditions determine *when* edge interpolation is warranted?

1.1 A Perceptual Problem

The existence of "gaps" in edges is an instance of the well known fact that information is lost when going from the distal stimulus to the retinal image. The perennial example of this is the loss of depth information: Scenes consist of a collection of objects in *three dimensions.* The retina is a *two-dimensional* surface; thus the projection of the scene onto the retina results in the loss of one dimension. A second example is the confounding, in the image, of surface reflectance and illumination. Reflectance and illumination are separate aspects of the scene that are conflated when projected into the image. The fact is, however, that we do not "perceive" our retinal images; rather, our perceptions correspond more closely to the properties of the underlying scene. We perceive in three dimensions, not two, and can discern both the intensity of an illuminant, and the reflectance of the illuminated surface. A problem of perception is to specify how these various scene properties are

surface is $x^2/80^2 + y^2/70^2 + z/100 = 0$. The illuminant direction is defined by the vector $(.2, -.95, .25)$. A lambertian reflectance function was used.

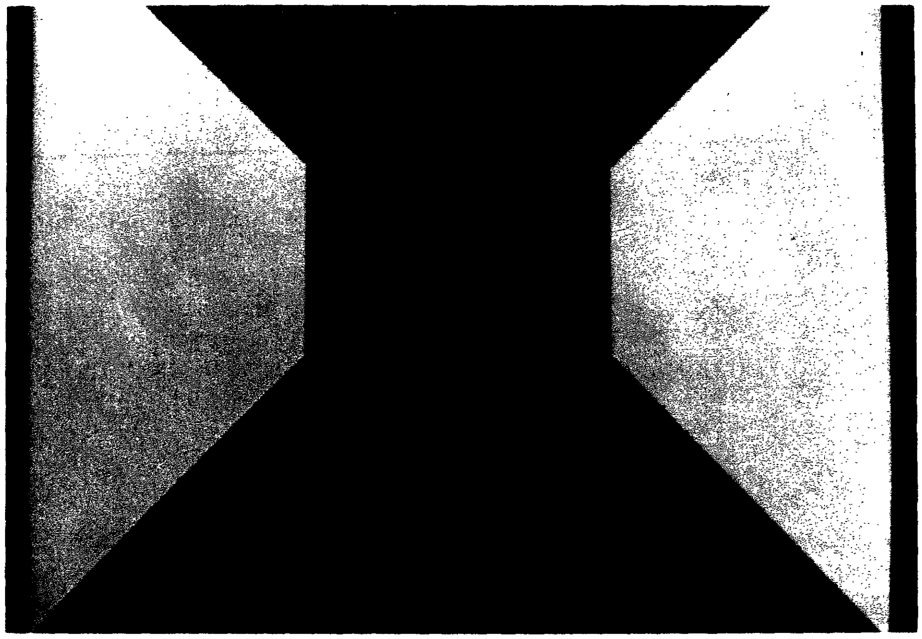

Figure 1.2: The zero contrast region clearly defined.

recovered based upon the often inadequate information given in the image. In particular, how are complete contours accurately perceived given that the corresponding edges appear in "patches along their length"?

A clue as to how the human visual system deals with the issue can be found in computational approaches to vision. Frequently, a *smoothness constraint* is invoked in order to solve some computational problem. The idea underlying smoothness constraints is that, on the average, scene properties do not vary substantially within a local neighbourhood. That is, if two adjacent points of a scene are compared, it is likely that the those points will be very similar to one another in numerous respects. The surface shape, reflectance, illuminance, and so forth will differ only sightly. This lack of change is mirrored in the image by a corresponding lack of change in image properties. Nevertheless, on occasion, adjacent regions of the image will differ substantially from one another, and these differences are informative, since they probably reflect a change in the scene. Edges themselves are an example of this: Edges are rapid changes in luminance which correspond to such things as occluding contours, shadow boundaries, and the like.

In any event, the utility of smoothness constraints is that gradual continuous change in the image implies continuity in some scene property; and, on the other hand, lack of continuity implies a change in the scene. With respect to the type of contour completion demonstrated in Figure 1.1, the reason why people see a contour here, where there is no edge, may have something to do with the manner in which the edges end on either side of the gap. If they gradually fade to zero contrast, that implies that the contours do not actually end. In the case of Figure 1.2, the gap is bounded by sudden changes in image intensity—it is surrounded by clearly defined edges and has a distinct parabolic shape. There is ample information in the image, in this case, to inform a perceiver that there really is a gap near the top of the figure.

This presence versus absence of clues at the edge endings is reminiscent of Grimson's (1981a, 1981b) work on surface interpolation. The computational problem in this case was to fit a three dimensional surface to a set of known disparity values. Grimson's solution can be summarized as follows. First, a pair of images is convolved using the Marr-Hildreth operator (Marr & Hildreth, 1980) and the zero-crossings are found. Second, the zero-crossings are matched between the two images and the disparities among them recorded. To this point, a representation of the local depth of the surface has been obtained, but only for those locations on the surface that gave rise to a zero-crossing. No disparity or depth values have been assigned to points between. How are the (correct) disparity values determined for these areas? Grimson (1981a) argued that the known values are used as anchors and a *smooth* surface is interpolated among them. The reason for preferring a smooth interpolation is this: Consider the changes in surface shape between a pair of matched zero-crossings. If surface depth is changing dramatically, then one would expect the image intensity values to reflect this change. That is, one would expect a third zero-crossing between the original pair. As there is in fact no such zero-crossing, then there is no significant change in the image intensities, and, hence, surface depth changed gradually and continuously in that region of the scene. Grimson referred to this state of affairs as the *surface consistency constraint* or, informally, as *no news is good news.* The lack of "news" in the image is evidence that changes in surface shape have been continuous in surface shape, and warrants a *smooth* interpolation of disparity values between the known values. On the other hand, a significant change in the image is evidence that smooth interpolation is not warranted.

1.2 An Hypothesis

A similar line of reasoning may hold with respect to edge endings and the kind of interpolation illustrated in Figure 1.1. When an edge ends suddenly, that state of affairs acts as a signal to the visual system and is indicative that something of note has occurred in the vicinity to cause the ending. Discontinuity in the contour is implied and the edge should be seen as ending. On the other hand, if an edge ends smoothly, continuity is implied, and the visual system ought to extend that edge until there is information to indicate otherwise. Such edge endings should lead to the perception of contours where there are no edges. In other words, unless there is some kind of "news" to indicate that the edge ends, the edge ought to be represented as continuing even where there is no luminance change in the image.

Consider Figures 1.3 and 1.4. In each figure, there is a set of four vertical bars each possessing a zero contrast region about their centres. The difference between the two figures is that the slope of the luminance gradient on either side of the gap is much greater in Figure 1.3 than in Figure 1.4. With respect to the left-most bar in Figure 1.3, the gap between the upper and

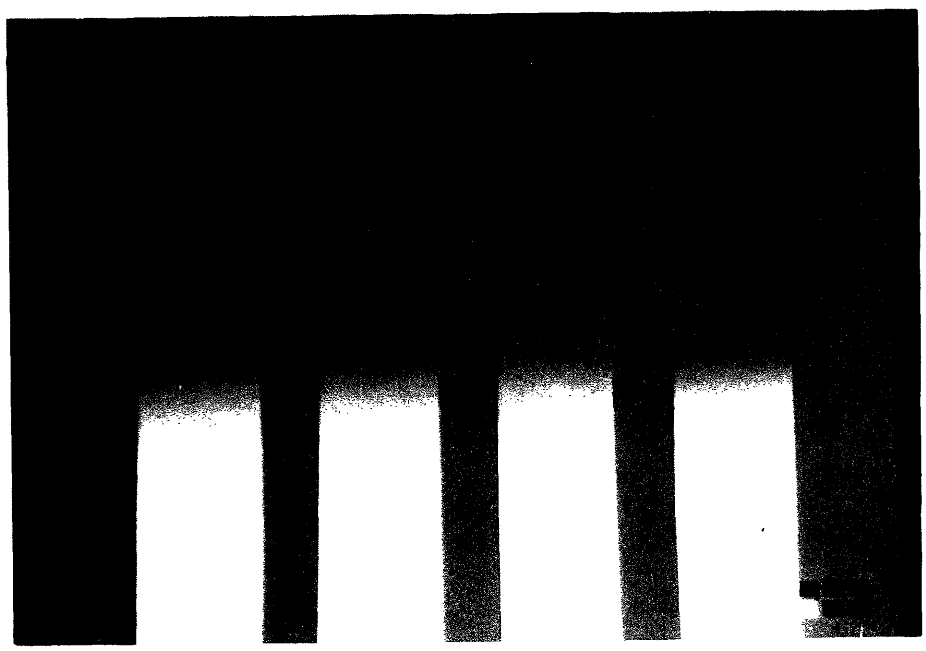

Figure 1.3: Edges ending abruptly.

lower portions of the bar is clearly visible and there is little or no tendency to experience a continuation of the vertical edges through the gap. By way of comparison, although the gap in the left-most bar in Figure 1.4, is detectable, there is an impression of a contour connecting the top and bottom portions of the bar. That is, the visual system is biased towards completing the vertical contours of the bar when the gap is bounded by smooth endings, and to not so complete when it is surrounded by more sudden endings.

The effect of other discontinuities on such completion is demonstrated by the remaining three bars in Figure 1.4. With respect to the second bar from the left, two straight edges have been added joining together the two pairs of vertical edges ending above and below the gap. The marking of the top and bottom of the gap in this way has the effect of making it more apparent; there is less of a tendency to perceive the vertical edges of the bar as extending across the gap. In the case of the third bar, the smoothness of the endings has been eliminated while retaining the luminance characteristics of the bar. Here, the abrupt endings on the left and right edges clearly inhibit any bias to perceive the bar as complete. Completely enclosing the top and bottom portions of the bar, as in the right-most bar in the figure, has a similar effect of making the gap clearly visible.

These examples lead to the following hypotheses regarding contour completion. First, smooth endings increase the likelihood of the perception of a contour not defined by a luminance change in the image. That is, smooth endings are likely to induce interpolation between them. Abrupt endings, in contradistinction, lead to the perception of "gaps". Second, smooth endings will lead to completion only under the condition that there is no other discontinuities restricting their continuation. Other edges provide clues that suggest further extension of the edge segment is not warranted. Thus, even when smooth endings are present in the image, they should not lead to interpolation if there is an event otherwise marking the gap between the endings—the gap should be clearly visible. Finally, this latter prediction has as a corollary that a stimulus which is in fact complete will be seen as gapped if there is information implying that there is a gap in the stimulus.[2]

[2]Obviously, this latter prediction will hold only when the edge is of low contrast; high contrast edges are detectable no matter what else is present in the image.

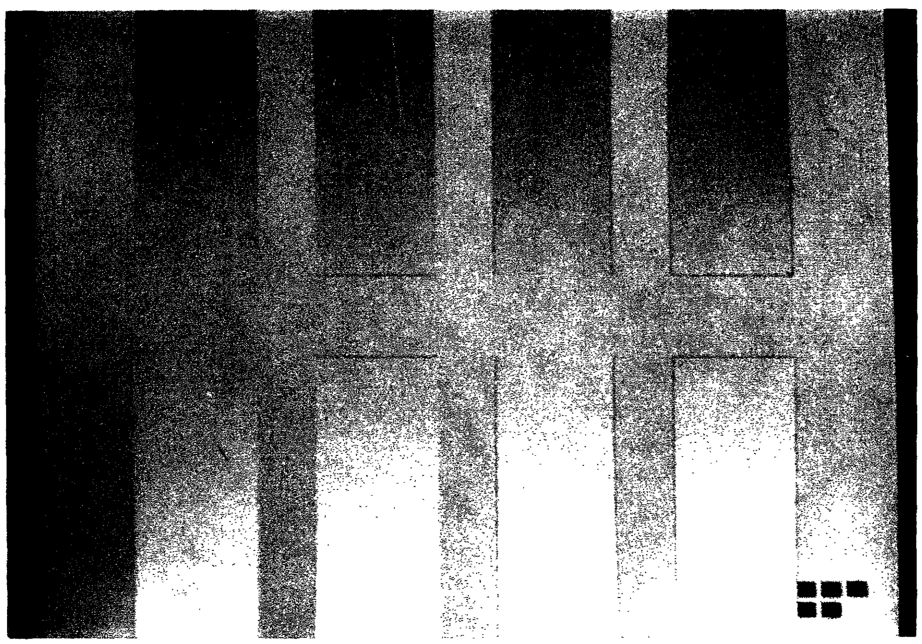

Figure 1.4: Edges ending smoothly.

1.3 Plan of the Thesis

These hypotheses were tested using two different paradigms and the results are reported, in detail, in the following chapters. They are described briefly here.

First, a detection task was used to determine whether subjects perceived contours when there was no luminance change in the stimulus. Subjects were presented with stimuli like those of Figures 1.4 and 1.3 and were required to judge whether there was a gap in the central region of the stimulus. The rationale behind this set of experiments was that any tendency to complete the contour should reflect in subjects' judgments—they should report "gaps" less often when the stimulus figure was designed to increase the likelihood of interpolation. The results confirmed the hypotheses. Stimuli involving smooth edge endings were judged "gapped" less often than those with abrupt endings. Also, "gapped" judgments were modified according to the presence or absence of various discontinuities. In general, the addition of any "news" increased the likelihood that a gap was perceived. In terms of ranking the effectiveness of the cues, a "box" drawn around the top and bottom of the stimulus induced the greatest number of "gapped" judgments, followed by abruptly ending vertical edges and, lastly, horizontal edges (see Figure 1.4). Furthermore, the presence of these cues caused the perception of a gap even when the stimulus was actually complete. This occurred when a low contrast but superthreshold edge was present in the "gapped region". These experiments are reported in Chapter 2.

Indirect evidence for completion is reported in Chapter 3. In a second set of experiments, the functional similarity between interpolated and real edges was assessed using the Poggendorff Illusion. To the extent that a gap causes an effect similar to that induced by real edges, here a Poggendorff Illusion, it can be concluded that the gap is not perceived as such; rather, at some level of processing within the visual system, the gap is filled in.

The results corroborated those reported in Chapter 2, although they were not as straightforward. To a first approximation, regions of physically zero contrast produced a significantly greater Poggendorff Illusion when bounded by smoothly ending edges as compared to those which ended abruptly. The gaps in the former stimuli did *not* behave like gaps, but as if a contour had been interpolated across the gap.

A clearly visible gap did not consistently result in a null effect. Instead, a misalignment effect was obtained in this condition, however, it was completely determined by the disparity between the angle of the gap and other aspects of the stimulus figure. The effect decreased as this difference decreased; and, when the difference was zero, no effects were observed. It appears that a gap bounded by abrupt endings did not induce a Poggendorff Illusion but induced a similar kind of effect based on the relationship between the shape of the gap and the rest of the stimulus figure.

The results of these two sets of studies are reported hereafter. In Chapter 4, the findings are discussed in the context of other related psychophysical phenomena including subjective contours, grating induction (McCourt, 1982), and phantom contours (Tynan & Sekuler, 1975). Also, the relationship between contour completion and edge detection algorithms is considered.

Chapter 2

Detection Experiments

The first set of experiments employed a gap detection task to test the hypothesis that smooth edge endings are likely to induce contour completion, and abrupt endings are not. Subjects were asked to view a number of stimulus figures and simply report whether a particular stimulus appeared to possess an area of zero contrast, or "gap", at its centre. Subjects were required to judge (a) if a single stimulus possessed a gap (Experiments 2.2 and 2.3), or (b) which of a pair of stimuli possessed a gap (Experiments 2.4 and 2.5).

The stimuli were varied in ways that, according to the hypotheses discussed in the previous chapter, increased or decreased the probability that subjects would perceive a gap. One of the variables studied was the slope of the luminance gradient on either side of a region of zero contrast. It was predicted that a gap would be seen more often when the gradient was relatively steep. Figures 2.1 and 2.2 illustrate this factor.

If contours are interpolated when continuity is implied by the ending, then if the continuity can be counteracted, interpolation should be checked. A second manipulation of the displays was the introduction of discontinuities whose purpose was to enhance gap detection. Faint edges, which are discontinuities in luminance, were used. These were placed near the gap in order to inhibit contour interpolation across the gap. These discontinuites constituted cues to the presence of the gap.

Four types of cues were used and these are illustrated in Figure 2.1. There is no cue marking the zero contrast area of the left-most stimulus, and it was predicted that this type of stimulus would be judged "gapped" the least. In this case, the edges ended smoothly on either side of the gap,

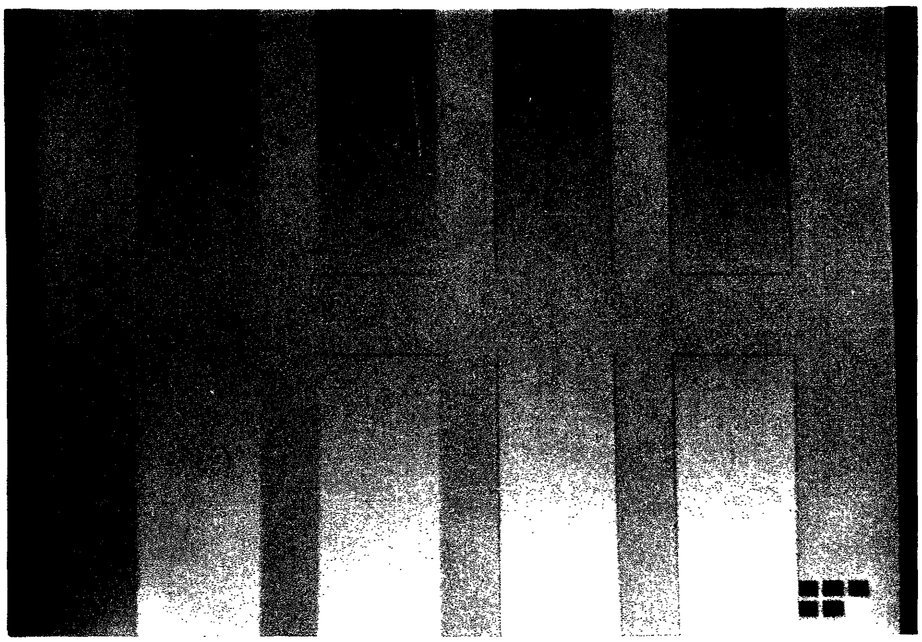

Figure 2.1: Examples of shallow gradients.

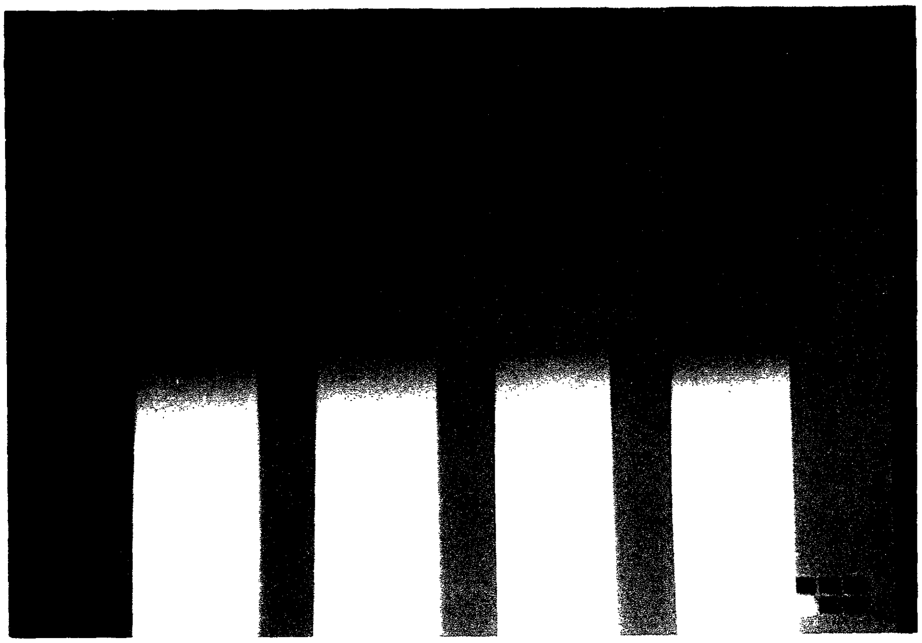

Figure 2.2: Examples of steep gradients.

and the gap was not emphasized in any way. The zero contrast area of the remaining three stimuli in the set was cued using three kinds of low contrast edges. Horizontal lines were used in one condition, joining the end points of the edge segments at the top and bottom of the zero contrast region. The second type of cue consisted of the placement of vertical lines that coincided with the vertical edges of the stimulus figure. These lines ended abruptly at the top and bottom of the gap. In this case, the edge segments did not end smoothly and it was predicted that this type of cue would also increase the perceptibility of the gap. The third cue consisted of a combination of the previous two, forming a "box" around the top and bottom of the stimulus, and it was expected that this too would enhance gap detection.

The hypothesis that cues enhance gap detection has as a corollary that a gap should be seen even when there is none in the stimulus. To test this, stimuli like those shown in Figure 2.3 were used to determine if subjects could be fooled into detecting non-existent gaps. The contrast of the vertical edges of these "complete" bars is low about their centres. It was predicted that since these edges were of low detectability, the cues would modify subjects' judgments and increase the frequency with which "gaps" were detected.

There was some concern that subjects would behave like a photometer given this kind of task. By asking subjects directly whether or not a particular stimulus has a gap, they are thereby informed that there may not actually be a gap in the stimulus and that, thus, they should adopt a fairly strict criterion. They will tend to look closely at each stimulus and examine it for any indication that it has a gap. This could lead to the uninteresting result that they could detect the gap almost 100 per cent of the time, and the results would say nothing about smooth vs. abrupt endings, nor the effects of the cues. In order to restrict such behaviour, therefore, stimuli were presented for a brief duration, and on occasion followed by a mask, so that the subjects could rely only upon their first impression of the stimulus.

Use of a brief presentation time gave rise to a concern that the low contrast edges of the complete stimuli (Figure 2.3) would be undetectable. If the contrast of these edges was subthreshold, then there is no reason to suspect that the cues account for any increase in the number of "gapped" judgments—the subjects would not see these edges even if cues were absent. Therefore, a group of subjects was run in a standard contrast threshold study in order to establish the minimum detectable contrast, and thereby insure that any results obtained with the complete stimuli were not due to the invisibility of

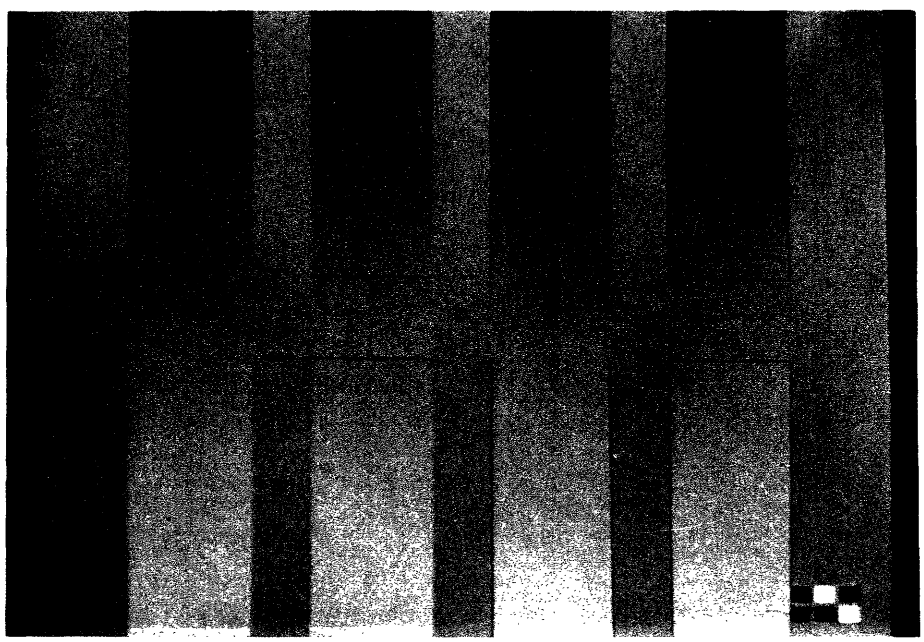

Figure 2.3: Examples of complete stimuli.

the edges in these figures. This control study is reported first.

Experiment 2.1

Method

Subjects. Eight University of Western Ontario undergraduate students took part in the study to fulfill a course requirement. Seven of the students were 19 years of age, the remaining one was 18. Five were male and three were female. All had normal or corrected to normal vision.

Apparatus and Stimuli. The following apparatus was used in all of the experiments. Stimuli were produced using a Matrox graphics system, consisting of two RGB-Graph/64-4 boards and a VAF-512/8 board. This system produced a display area 512 pixels wide by 480 pixels high. Given the viewing distance of 98.7cm, each pixel subtended 1.69′ of visual angle horizontally by 1.35′ vertically. The stimuli generated by this system were displayed on an Electrohome RGB monitor, model number 38-D013101-60. This apparatus was capable of producing 256 grey levels, ranging from 4.3 candelas ("black") to 233 candelas ("white"). A Volker-Craig VC4152 terminal was used to record subjects' responses. The experiments were computer run on a National Semiconductor 16032 minicomputer under the Unix operating system. A chin rest was used to stabilize subjects' heads.

The stimuli consisted of a single vertical bar 10.82 deg of visual angle high by 1.97 deg wide, displayed upon a larger background of constant luminance (97 candelas). The dimensions of the background were 10.82 deg high by 14.42 deg wide. The bars were presented 1 deg to the right of a central fixation cross. The reason for this was that, in Experiments 2.4 and 2.5, a pair of bars were presented to the left and right side of fixation. It was desired that the detection thresholds estimated in the current study be obtained under similar conditions, in which sensitivity is potentially decreased by an off centre display.

The contrast of the bar varied from trial to trial. Contrast was calculated using the formula $(L_{background} - L_{bar})/(L_{background} + L_{bar})$. Contrast ranged from zero per cent to a maximum of 4.41% darker than the background. On each trial, the stimulus bar was presented for a duration of 864 msec; an abrupt stimulus onset was used.

A dual staircase procedure, based on Carterette (1984), was used to determine subjects' contrast thresholds. For one of the staircases, the stimulus contrast was set initially to zero per cent, and, for the other, to the maximum contrast. Step size was a constant 1.5 candela increase or decrease of the bar's luminance (the resolution of the equipment). Criterion was set at eight turnarounds, where the last six of these were used to estimate the threshold. Wetherill's decision rule two (Carterette, 1984) was used to reverse the direction of the staircase: Contrast was lowered only after subjects had indicated that they had seen the stimulus for two successive presentations (within the same staircase); contrast was increased after a single failure to detect the bar. This procedure estimates the 70.71% threshold.

Procedure. In order to reduce reflection from the monitor screen, the room lights were turned off at the beginning of the experimental session. Under these conditions, the room was not completely dark, only dim. While subjects were dark adapting, the procedure was explained to them. They were instructed to place their heads in the chin rest and to gaze at the fixation cross on the monitor screen. They were told that they would be presented with a number of vertical bars, slightly to the right of fixation, and that they were to indicate whether they had seen the bar or not. If they had, they were to press the "y" key on the terminal keyboard, the "n" key otherwise.

The sequence of events was then explained. At the beginning of each trial, the screen was set to the background luminance. During this time, the computer generated the display specific to the trial, but did not show it. The terminal beeped to indicate to the subject that the stimulus was ready for presentation. Subjects were told that they were to prepare themselves for its presentation, gaze at the fixation cross, and to press the space bar on the keyboard when they were ready. Shortly after pressing the space bar, the stimulus bar was presented and then erased. The appearance of the stimuli was of a bar flashed briefly against a constant grey background. The terminal beeped a second time to signal that the bar had been shown. Subjects were instructed to press the appropriate key at this point, indicating whether they had seen a bar during the interval between the first and second beep. After they had made their response, the computer proceeded to the next trial and the process was reiterated until the staircase criteria were met. The interval between trials was approximately 10 seconds.

Due to equipment failure, the stimulus was presented, on occasion, for longer than the desired interval, and remained on the screen for about four seconds. This was an infrequent and random occurrence (once or twice during an experimental session). Subjects were instructed that whenever this happened, they were to press the "r" key on the keyboard, and the trial was repeated at later time. After all trials were completed, the subjects were given a short written account explaining the purpose of the study, and any questions that they had about the study were answered.

Results

For each subject, the thresholds estimated from the ascending and descending staircase were combined and a mean computed. The contrast threshold for each of the eight subjects was 0.35%, 0.79%, 0.4%, 0.85%, 0.4%, 1.24%, 0.52%, and 0.4%. The mean of these estimates is 0.62%. Given the luminance of the background used in the current experiment, the minimum contrast the apparatus was capable of generating was 0.7%, a value slightly higher than the obtained mean threshold. In other words, on average, the subjects could detect an edge possessing the minimum possible contrast.

The reason that the threshold is less than the minimum possible contrast is a function of the manner in which the threshold is calculated. The threshold is estimated as situated somewhere between the minimum level of contrast at which subjects "always" detect the bar, and the maximum level at which they "never" detect the stimulus. In this particular study, the two levels in question were the minimum possible contrast and zero contrast.

Experiment 2.2

The next two experiments measured the absolute detectability of gaps. Subjects were shown single bars and were asked to say whether the bar possessed a gap. Later studies measured relative detectability by presenting subjects with pairs of bars.

Method

Subjects. Ten University of Western Ontario undergraduate summer school students were paid $10 for their participation in the study. Their ages ranged from 19 to 29. Seven were female and three were male. All had normal or corrected to normal vision. All were naive with respect to the purpose of the study.

Stimuli. There were ten displays each consisting of a single vertical bar centred on the background. The luminance of the stimuli, and hence the contrast between them and the background, varied as a ramp function of the bar's height. Each bar was darker than the background at the top of the display, and became progressively lighter towards the bottom of the display; bar luminance equalled the background luminance at the centre of the display. Five bars served as experimental stimuli and possessed a region of zero contrast about their centres, 1.35 deg of visual angle in height. That is, their luminance was constant and equal in value to that of the background beginning at a distance of 4.74 deg from the top of the display and ending at 6.09 deg. From here on, the luminance of the bar increased. These stimuli are referred to as "gapped".

The luminance of the five control stimuli did not likewise plateau, but continuously increased from the top to the bottom of the display. As such, these stimuli did possess a gap, albeit much smaller than that of the gapped bars (physically, one pixel or $1.35'$ in size). These latter five bars are termed "complete". The luminance profiles of the complete and gapped bars are illustrated in Figures 2.4 and 2.5, respectively. As can be seen in the figures, if the maximum ("white") or minimum ("black") luminance the apparatus was capable of generating was reached, the bar's luminance was set to "white" or "black" over its remaining length. The steepest gradient depicted in Figure 2.5 was not used in the present study.

The bars varied in terms of the slope of their luminance profiles, which were approximately linear. The slope of the profile is reported here as:

$$\text{gradient slope} = \Delta L / \Delta s$$

where ΔL is the change in luminance (in cd/m^2) and Δs is the change in position along the height of the bar (in degrees of visual angle). Larger slopes denote steeper luminance gradients. For example, a step change in luminance (the steepest gradient possible) has a value of infinity, while no gradient is denoted by "zero"; the latter represents a situation in which there is no stimulus bar in the display. The slopes of the luminance gradients of the five complete bars were 1.96, 3.95, 9.45, 18.23, and 32.20 cd m^{-2} deg^{-1}. The comparable slopes for the five gapped bars were 2.45, 4.47, 10.50, 18.72,

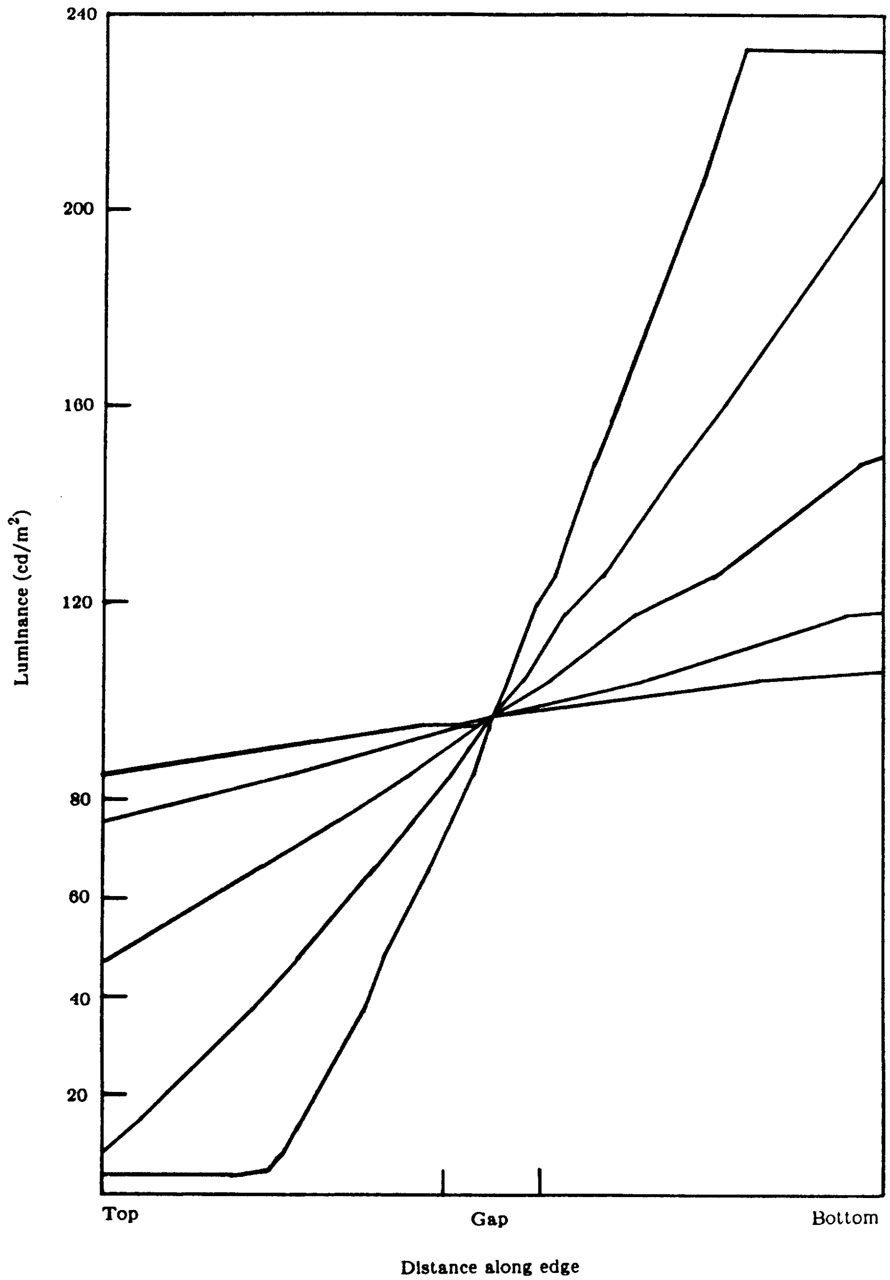

Figure 2.4: Luminance profiles of the complete stimuli.

and 31.87 cd m^{-2} deg^{-1}. The luminance gradients of the complete bars were selected such that their overall luminance matched that of their gapped counterpart. The luminance of the bars ranged from 85–104 candelas, 76–118 candelas, 47–148 candelas, 8–204 candelas, and 4.3–233 candelas, respectively, for each of the five pairs of complete and gapped bars.

Note that the five levels of gradient slope effectively changed the overall contrast of the stimulus bars. Bars with steeper gradients were of greater average contrast than those defined by shallower gradients. In particular, that area of the complete stimuli, corresponding to the zero contrast region of the gapped stimuli, varied in its average contrast. The mean contrast of this area was 0.88%, 1.05%, 1.40%, 4.90%, and 8.10%, respectively, for the five complete bars. Figure 2.3 shows the 1.05% contrast stimuli.

A second feature of the displays was the type of cue marking the top and bottom of the central region. Four cues were tested: (1) a no cue condition in which the stimulus bar appeared as described above, (2) a condition in which a pair of horizontal lines were placed at the top and bottom of the central region, connecting the left and right vertical edges of the bar, (3) four vertical lines collinear with the outer edges of the bar, each ending abruptly at the top and bottom of the central area, and (4) a combination of horizontal and vertical lines which formed a "box" around the top and bottom sections of the bar (see Figures 2.1 and 2.3). The cues were placed at the top and bottom of the gap in the case of the gapped bars, and, for the complete bars, at the same spatial locations as on the gapped bars.

The luminance of the cues was set so that they were slightly darker than the bar at the top and bottom of the central region. A constant decrease of 18 candelas was used. For the gapped stimuli, this meant that the luminance of the cue lines was consistently 78.57 candelas, since the luminance at the top and bottom of the gap was equal to the background. For the complete bars, the luminance of these lines varied since, here, the luminance at the top and bottom of the central region of the bar differed for each of the five stimuli. The luminance of the cues at the top were 76.86, 75.84, 72.42, 66.14, and 56.64 candelas, respectively. At the bottom, they were 80.55, 81.30, 85.40, 92.91, and 103.16 candelas, respectively.

All 40 displays were shown to each subject. Subjects were run in two blocks of 200 trials each; within each block, each display was presented on five separate occasions. Thus, each subject made a total of 10 judgments

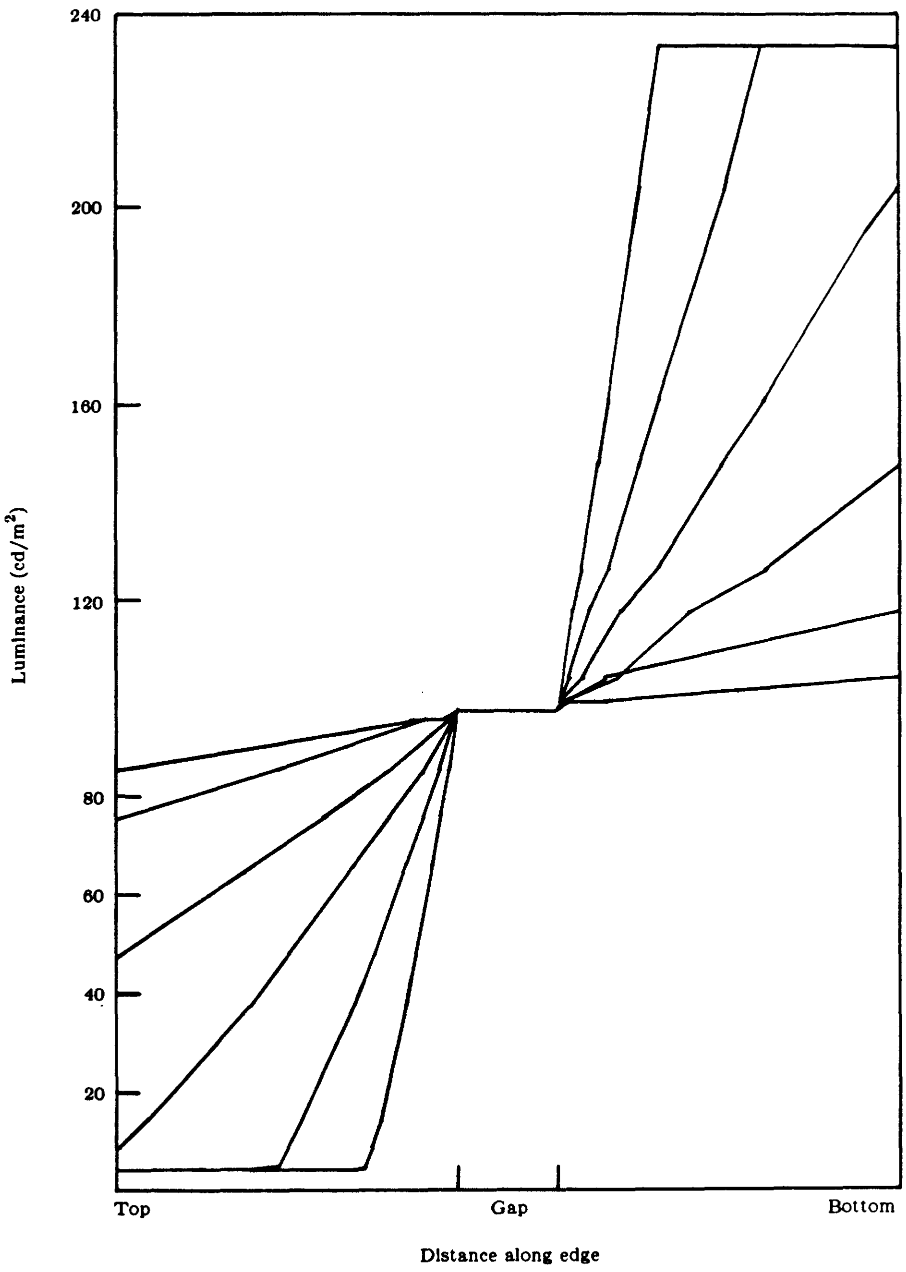

Figure 2.5: Luminance profiles of the gapped stimuli.

per display. The order of presentation of the various displays was randomized for each subject and for each of the two blocks of trials. Each display was presented for a duration of 864 msec followed by a "white" mask (233 candelas); an abrupt stimulus onset was used. Seven displays were used for demonstration purposes; these consisted of some of the displays described above, as well as others of similar construction.

Procedure. Lighting conditions were dim, as before. Subjects were shown the first example display, an no-cue gapped bar with a contrast gradient of 31.88 cd $m^{-2}\,deg^{-1}$. The gap was readily visible in this display. Subjects were told that they were going to view many similar figures and that their task was to judge whether the bar had a gap in its vertical centre. It was emphasized that if the bar did possess a gap, then that gap would be always centrally located and equivalent in size to that shown in the example. The second example display was then shown, a no-cue gapped bar with a gradient of 4.47 cd $m^{-2}\,deg^{-1}$. Subjects were told that the gap was not always as easily detected as that in the first display. This second display illustrated that point. The third example display was then shown, which consisted of a no-cue complete bar with a gradient of 9.45 cd $m^{-2}\,deg^{-1}$. This display was used to exhibit the appearance of a complete bar with a shallow luminance gradient, that sometimes the vertical edges in the central region would be difficult to see even though there was no gap in the stimulus. Finally, a gapped bar with a gradient of 10.49 cd $m^{-2}\,deg^{-1}$ and the horizontal lines cue was shown. Here subjects were told that, in addition to the bars, there frequently would be various faint lines near the central area of the bar. It was emphasized that these lines would be displayed regardless of whether the bar had a gap or not, and that the lines might help or might hinder them in their task. They were instructed to decide whether the stimulus did indeed possess a gap, to the best of their ability. During this phase of the experiment, the displays were not flashed briefly, but were shown for longer durations so that the subjects could inspect the figures and see them clearly.

The mechanics of the experiment were then explained to the subjects. At the beginning of each trial, the monitor screen was black and remained so while the computer generated the next display to be shown. When the display was ready, the terminal beeped as a signal to the subjects. Subjects were told that when they were prepared to view the display, they were to press the space bar on the terminal any time after the beep. They were advised that it was best to look approximately at the centre of the monitor screen before

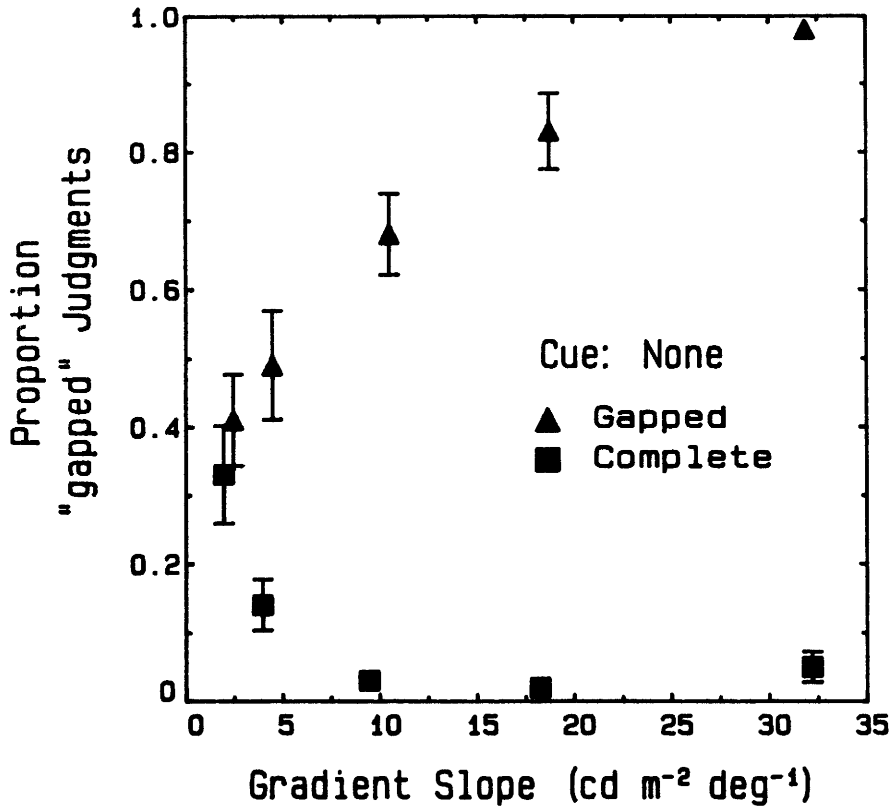

Figure 2.6: Proportion "gapped" judgments as a function of gradient slope, no cue (Experiment 2.2).

pressing the space bar (no fixation cross was provided in the current study). Shortly after pressing the space bar, the display was presented, followed by the mask. Subjects were instructed to press the "y" key on the terminal if they thought the stimulus bar did possess a gap, or the "n" key if they felt it did not possess a gap. They were told to guess if necessary. After pressing one of these keys, the monitor screen went blank, and the procedure reiterated.

The subjects were then run in a block of eight practice trials in order to familiarize them with the equipment and procedure. Following the practice trials, the subjects were run through the first block of experimental trials. There then followed a short break of approximately five minutes followed by the second block of experimental trials. As in the previous experiment, the stimulus occasionally was presented for a longer period than desired. The trials on which this happened were noted, and these trials were repeated at the end of the block in which they occurred. After all trials were completed, the subjects were suitably debriefed.

Results

The proportion of "gapped" judgments made by each subject for each display was calculated. The mean proportion "gapped" judgments collapsed over all subjects are shown in Figures 2.6–2.9. The slope of the contrast gradient is given on the x-axis and the proportion "gapped" judgments on the y-axis. Triangles represent judgments made in conjunction with bars that were truly gapped; squares represent responses made when presented with complete bars.

For gapped bars, it is obvious that the steeper the gradient, the more likely subjects were to judge the bar as possessing a gap. Subjects' performance was near perfect for the steepest gradient, and inspection of Figures 2.6 through 2.9 reveals than the type of cue had little effect in this case. The gap was clearly detectable when the contrast gradient was steepest.

For shallower luminance profiles, subjects were less inclined to make a "gapped" judgment, and here their judgments were influenced by the type of cue surrounding the gap. They were least likely to judge a stimulus as gapped if there was no cue to its existence, and increasingly more likely to so judge as the type of cue changed from horizontal lines, to vertical lines, to a box.

Since the means and variances of this data are not independent, individual subjects' proportions were transformed according to the arcsin transformation (Kirk, 1968), and these transformed scores were analyzed according to a five (gradients) by four (cues) analysis of variance with repeated measures on both factors. There was a main effect of gradient slope, $F(4, 36) = 24.02, p < .001$, and a main effect of cue, $F(3, 27) = 15.87, p < .001$. There was also a significant interaction between the slope of the gradient and the type of cue, $F(12, 108) = 4.63, p < .001$. The interaction indicates that the zero contrast region was most apparent when the luminance gradient was steep, and that subjects could easily perceive it in this condition regardless of any enhancement provided by the cues. Indeed, the cue was superfluous when the gap was already so clearly defined. When the gap was not as distinct, the type of cue was a factor in terms of increasing the number of "gapped"

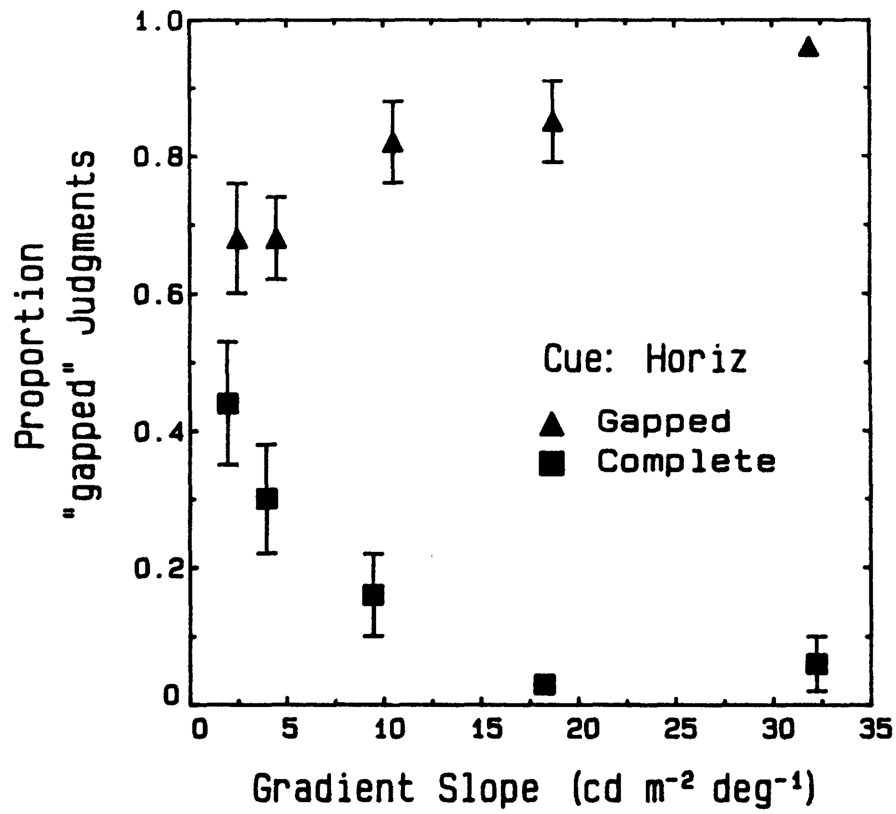

Figure 2.7: Proportion "gapped" judgments as a function of gradient slope, horizontal lines (Experiment 2.2).

judgments. Different cues were more influential than others.

A complementary pattern of results was obtained with complete bars. As can be seen in Figures 2.6–2.9, the proportion "gapped" judgments became less as the slope of the luminance gradient increased. The subjects were virtually perfect when the gradient was steep, and made more errors as the slope became less steep. Inspection of Figures 2.6 through 2.9 reveals that, for shallow contrast gradients, the proportion of "gapped" judgments increases as the cue went from none, to horizontal lines, to vertical lines, to the box, increasing by as much as 29%.

Proportions were again transformed and analyzed according to a five by four repeated measures analysis of variance. The results paralleled those obtained using gapped stimuli. Both main effects were significant, ($F(4,36) = 69.82, p < .001$, and $F(3,27) = 12.74, p < .001$, respectively) as well as the

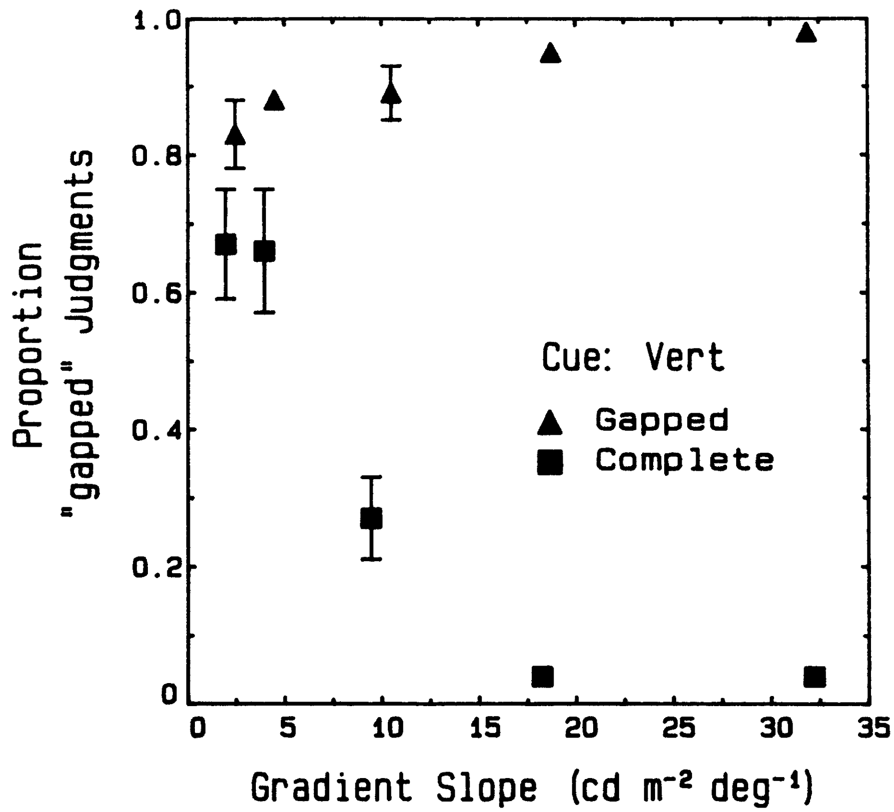

Figure 2.8: Proportion "gapped" judgments as a function of gradient slope, vertical lines (Experiment 2.2).

interaction, $F(12, 108) = 5.75, p < .001$.

The results are summarized as follows: First, when the zero contrast region was clearly visible, or alternatively clearly absent, then the subjects perceived the stimuli accurately and made their judgments in accordance with their perception. In this case, the presence of various cues had no influence on the perception of the central region. Second, when the central region was not clearly defined in terms of whether it possessed an actual gap or a faint edge, then subjects made more errors, judging gapped stimuli as complete and complete bars as gapped. Third, their judgments for the shallow gradient stimuli were influenced by the type of cue. It appears that there being no cue in the display was most conducive to incurring a "complete" judgment. The addition of any cue increased the proportion of "gapped" judgments. For the different cues, horizontal lines had a lesser effect than

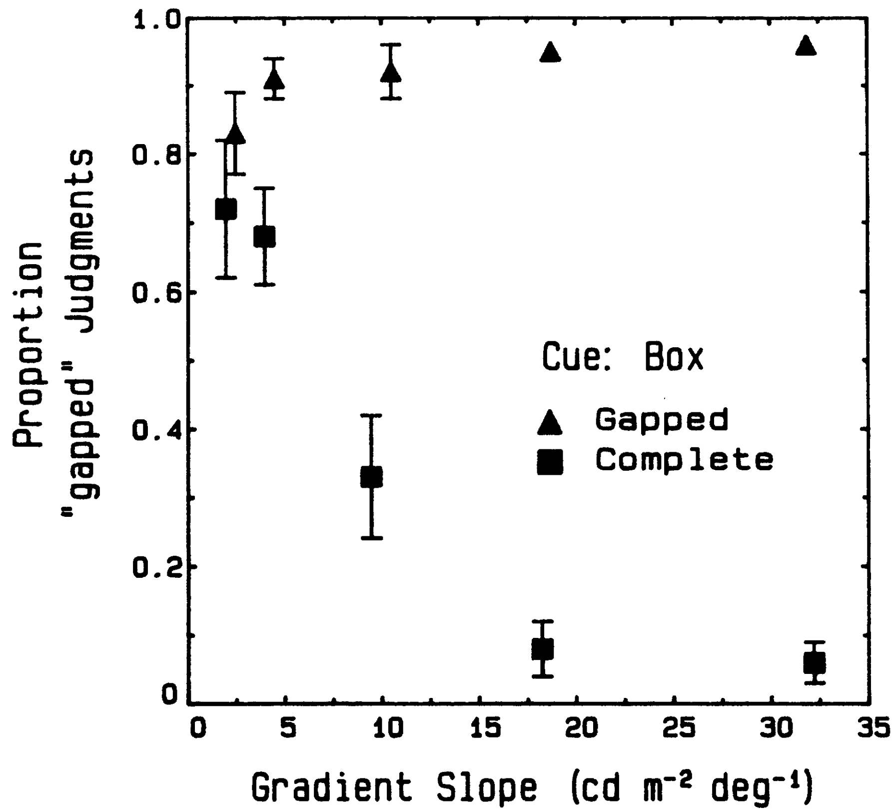

Figure 2.9: Proportion "gapped" judgments as a function of gradient slope, box (Experiment 2.2).

either vertical lines or a box. The latter two cues induced approximately equal proportions of "gapped" judgments. These results indicate that when edges are difficult to see, subjects relied on other information in the stimulus to determine whether there really was a gap in the bar.

Experiment 2.3

Experiment 2.3 replicated the previous study using bars defined by *reversing* gradients. That is, the bars were defined by gradients of increasing luminance from the top of the display to the centre, then decreasing luminance over the rest of their length. The luminance of the bars in the previous experiment increased over their entire length (see Figures 2.4 and 2.5). The result of designing the stimuli in this manner was that the contrast of the central region of the complete stimuli was constant over its length.

Method

Subjects. Ten first year undergraduate University of Western Ontario students participated in the study to fulfill a course requirement. Their ages ranged from 19 to 27. Seven were female and three were male. All had normal or corrected to normal vision. None had participated in previous experiments.

Stimuli. The stimuli were similar to those of the previous experiment, with the following differences. Three bars were used, whose luminance profiles were again a ramp function of the height of the bar. The slope of the gradient was equal to 4.47 cd $m^{-2}\,deg^{-1}$ for all three bars; however, in the present experiment, the direction of the gradient reversed at the centre of the bar. That is, as before, the bars were darker at the top of the screen and became progressively lighter in luminance towards the centre of the display. The luminance of the bar did not continue to increase upon reaching the bottom of the central region, but decreased in luminance towards the bottom of the screen. In short, the intensity profile of the bottom half of the bars was a mirror image of their top halves.

The luminance of a region 1.35 deg about the centre of the bars was held constant, and the three bars differed with respect to the contrast of this region. In one case, the luminance of the central region was set to that of the background, resulting in a gap of zero contrast. For the second bar, the contrast between this region and the background was set to 0.88%, and the third was 2.8%. These contrasts were obtained by decreasing the intensity of the first bar everywhere by a constant amount. There was thus an gapped bar, a complete bar with a relatively faint edge through its centre, and a complete bar with a stronger edge. The three luminance profiles were crossed with the four types of cues. The cue lines were a constant 18 candelas darker than the luminance of the bar at the top and bottom of the central region. An example is shown in Figure 2.10; here the gapped bar.

Each of the 12 displays was presented on 10 occasions resulting in a block of 120 trials. Each subject received a different random order of the trials. In the current study, stimulus onset was abrupt, preceded and followed by an blank screen equal in luminance to that of the background.

Procedure. The procedure was the same as that used in Experiment

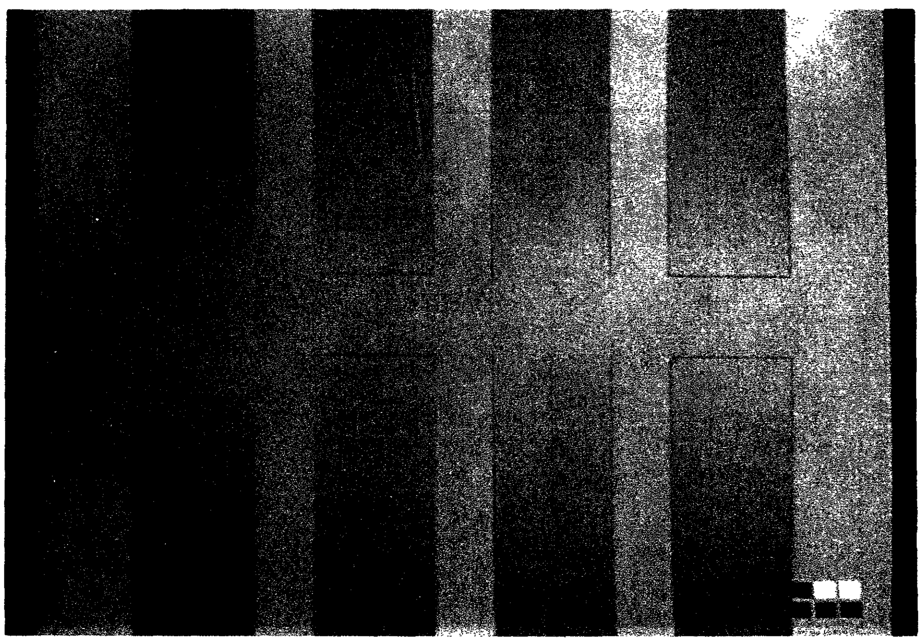

Figure 2.10: Example of displays used in Experiment 2.3.

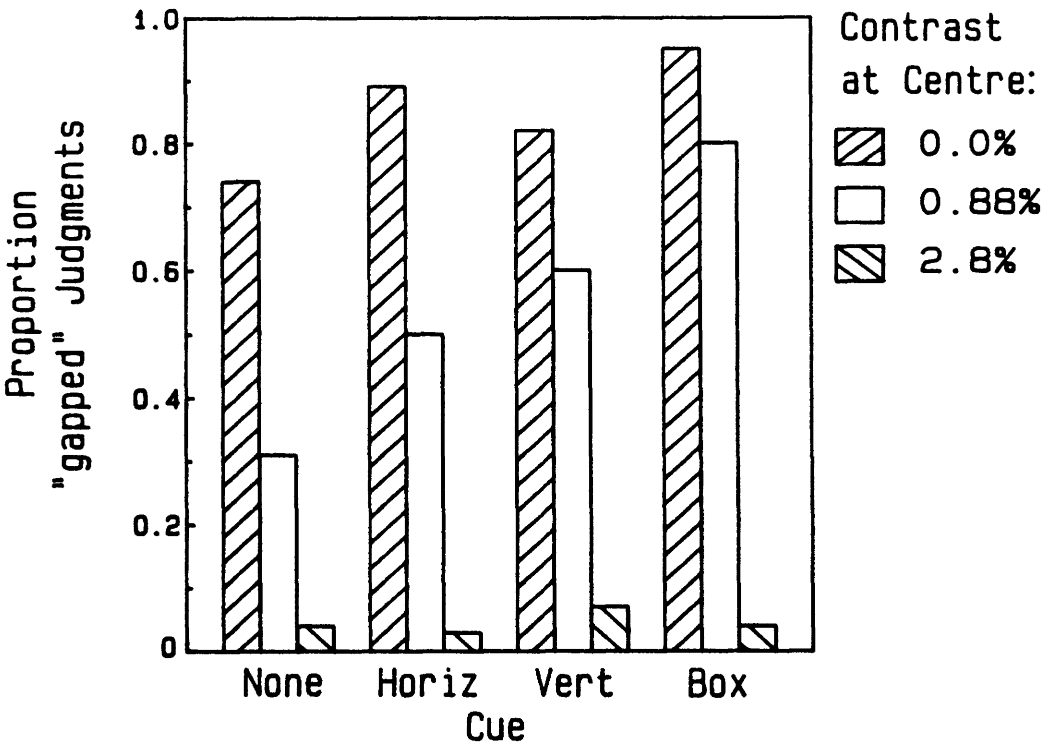

Figure 2.11: Proportion "gapped" judgments as a function of cue type (Experiment 2.3).

2.2, with the following differences. A different set of demonstration/practice displays was used in order to match the reversing luminance profiles of the displays used in the experiment proper.

The sequence of an individual trial differed slightly from that in Experiment 2.2. At the beginning of the trial, before the terminal beeped, monitor screen luminance was identical to that of the background of the displays. When the subject initiated a trial, this field was replaced by the particular stimulus for that trial, which was subsequently erased. The appearance of the stimuli was that of a bar flashed briefly against a grey background.

Results

The proportion of "gapped" judgments made by each subject for each display was calculated. The means are shown in Figure 2.11. The type of cue is labelled on the x-axis, and the proportion "gapped" judgments on the y-axis.

The histogram filled with right diagonal lines shows the results obtained with the gapped stimulus; the open histogram, the results obtained in conjunction with the bar completed using a relatively faint edge; and the left diagonal lines histogram, the bar completed using a higher contrast edge.

The proportion of "gapped" judgments an individual subject made for each contrast by cue condition was transformed according to the arcsin transformation and these transformed scores were analyzed according to a three (contrast of central region) by four (type of cue) analysis of variance. Both main effects were significant; with respect to contrast: $F(2,18) = 174.70, p < .001$, and with respect to cue: $F(3,27) = 8.53, p < .001$. The interaction between contrast and type of cue was also significant, $F(6,54) = 6.51, p < .001$.

These effects can be seen in Figure 2.11. As in the previous experiment, subjects rarely judged the bar as "gapped" when the edge in the central region of the bar was clearly visible (the histogram filled with left diagonals in Figure 2.11). The type of cue also failed to influence subjects' decisions in this situation. But, when the edge in the central region was of lower contrast, the type of cue did modify subjects' judgments. The addition of a cue caused an increase in "gapped" responses as compared to the no cue condition. vertical lines and the box had a greater effect than did horizontal lines.

The results obtained with a truly gapped bar were similar to that observed in the previous experiment, although there was a tendency to judge the gapped bar as "gapped" more often in the present experiment. The gapped bar here was judged "gapped" 74%, 89%, 82%, and 95% of the time, for each of the cues. In the previous experiment, the gapped bar that had the same contrast gradient was judged "gapped" 49%, 68%, 88%, and 91%, respectively. When no cue was present, the gap was detected more often in the current study as compared to Experiment 2.2. Since a similar sort of difference was obtained between Experiments 2.4 and 2.5, discussion of these results is postponed until the general discussion.

Experiment 2.4

An implication of the previous two experiments is that some stimuli looked "more gapped" than others. Consider a gapped bar defined by a shallow gradient without any cue lines. Such a stimulus was judged "gapped" 49% of the time in Experiment 2.2. By comparison, the addition of horizontal lines above and below the gap increased the proportion of "gapped" judgments to 68%, for the same stimulus bar. Furthermore, a physically complete bar defined by a comparable luminance profile was judged "gapped" 66% of the time when the vertical cue lines were placed on the stimulus. In other words, a complete stimulus was judged "gapped" more often that a similar no-cue *gapped* stimulus. These results suggest that the cues could be ranked according to the degree with which they made the stimuli appear "gapped". The next two experiments were an attempt to construct such a ranking.

The method of paired comparisons (Guilford, 1954, chapter 7) was used to construct a unidimensional scale of how "gapped" the different cues made the bar appear. Where might the stimuli fall on such a scale? Consider the second steepest gradients tested in Experiment 2.2, for both the gapped and complete bars (4.46 and 3.94 cd $m^{-2}\,deg^{-1}$, respectively). Inspection of Figures 2.6–2.9 reveals the following ordering of these eight stimuli ("g" denotes the gapped bar, "c" the complete; the no-cue condition is denoted by "0", horizontal lines by "h", vertical by "v", and the box by "b"): c0, ch, g0, cv, cb=gh, gv, gb. Overall, the complete stimuli were judged "gapped" less often than the stimuli possessing a gap but there was some overlap. Also, certain cues increased the proportion of "gapped" judgments more than others. The order was no-cue, horizontal, vertical, and finally box, irrespective of whether the bar was gapped or not. It was predicted that the scale found would be qualitatively similar to the above ordering of the stimuli.

Method

Subjects. Ten University of Western Ontario introductory Psychology students participated to fulfill a course requirement. All had normal or corrected to normal vision, and none had been in either of the previous experiments. Their ages ranged from 17 to 19 years. Eight were female and two were male.

Stimuli. Eight bar stimuli were appropriated from those described in Experiment 2.2. Four of these were gapped and four were complete. The slope of the luminance gradient for the gapped stimuli was 4.46 cd $m^{-2}\,deg^{-1}$, and that of the complete bars was 3.94 cd $m^{-2}\,deg^{-1}$. Within each set of four, the central region was marked by one of the four cues.

In the present study, a display consisted of a pair of these bars separated by 2 deg about a central fixation cross. Every possible pair was used with the exception of pairing a stimulus with itself, resulting in a total of 28 pairs. Two sets of displays were made, one set of 28 in which one of the bars in each pair was presented on the left side of the display, and one set of 28 in which it was presented on the right side. Each of these 56 displays was presented five times. Each pair therefore was presented 10 times giving a total of 280 trials. Subjects were run in two blocks of trials, first in a block of 168 trials and then in a block of 112 trials. A different random order of presentation was used for each block, and for each subject. There was a five minute break between blocks. As in Experiment 2.3, no mask was used. A set of example displays were constructed and used to describe the task to the subjects, and to give them practice. These displays also consisted of pairs of bars.

Procedure. The procedure was the same as that employed in Experiments 2.3 with one difference. The judgment in the present study concerned a pair of bars—subjects were instructed to choose which of the pair had the gap. They were told that this decision would be fairly difficult on occasion, and that they should then ask themselves which of the two bars appeared "more gapped". Regardless of how difficult the comparison was, they had to choose one of the bars. If they thought the bar to the left of centre was the gapped bar, they were to push the "z" key on the terminal, and, if they chose the right hand bar, the "/" key was to be pushed.

Results

The mean proportion "gapped" judgments for each stimulus in a given pair was computed across all subjects. The mean proportions are shown in Table 2.1. The column and row headings designate the various stimulus bars both with respect to whether it was complete and in terms of the cue associated with that stimulus. Gapped bars are denoted by "g", and complete bars by "c". The no-cue condition is given by "0", horizontal lines by "h", vertical lines by "v", and the box by "b". Each cell in the table designates the frequency with which the column stimulus was judged "gapped" when compared to the row stimulus. For example, to determine the proportion of "gapped" judgments the gapped bar/no-cue stimulus (g0) received when paired with an gapped bar/horizontal lines stimulus (gh), read down the column labelled "g0" to the row labelled "gh". The .45 found there indicates that subjects judged the no-cue stimulus as having a gap 45% of the time when compared to the one involving horizontal lines (alternatively, the bar

Table 2.1:

Proportion matrix for paired comparisons (Experiment 2.4).

Bar	c0	ch	g0	cb	cv	gh	gv	gb
c0	–	.50	.76	.74	.74	.71	.80	.82
ch	.50	–	.66	.70	.66	.79	.79	.86
g0	.24	.34	–	.55	.47	.55	.57	.65
cb	.26	.30	.45	–	.50	.51	.67	.67
cv	.26	.34	.53	.50	–	.41	.61	.68
gh	.29	.21	.45	.49	.59	–	.62	.65
gv	.20	.21	.43	.33	.39	.38	–	.54
gb	.18	.14	.35	.33	.32	.35	.46	–
Scale	0.0	0.0326	0.5854	0.5863	0.5953	0.6118	0.8811	1.0093

with horizontal lines bar was judged "gapped" 55% of the time, for the same comparison).

The proportions were used to generate the scale. Each proportion in the table is assumed to correspond to an area under the normal distribution, and is converted to the corresponding z-score. For example, a proportion of .50 is converted to a z of zero. The columns of the resulting table of z-scores are summed and a mean is calculated. The lowest mean is set to zero, and the other means are adjusted appropriately. The mean z-scores calculated in this way represent the relative positioning of the stimuli on a unidimensional scale; in this case it is scale of the degree to which a particular bar appears "gapped". The scale values for each stimulus are shown in the bottom row of Table 2.1.

The scale is shown graphically in Figure 2.12. Because some of the stimuli were very close, the scale has been drawn separately for the complete (top line) and gapped bars (bottom line). There is only one scale here; it is drawn twice only to increase the legibility of the figure.

For complete bars, the ordering was not exactly as expected, although the deviation was not that severe. The no-cue stimulus was judged the least gapped, followed by the bar marked with horizontal lines, the box, and finally the vertical lines. The only difference between this ordering and that

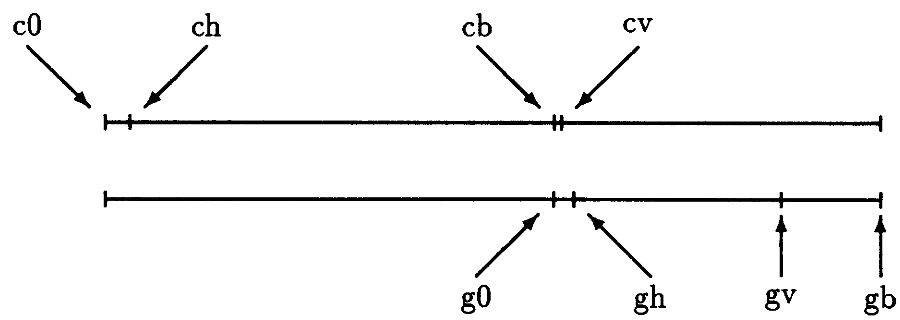

Figure 2.12: The scale for the stimuli in Experiment 2.4.

expected was that the vertical lines and box cues were switched, although they were nearly equal. Note that they did have higher scale values as compared to the no-cue and horizontal cue stimuli.

For gapped bars, the ordering was as expected—the no-cue bar was judged "least gapped" when compared to the rest. Horizontal markings were judged as appearing slightly more "gapped", while vertical lines and the box were clearly the most "gapped" in appearance. In addition, there was some overlap in the scale values of the gapped and complete stimuli, as predicted. Figure 2.12 shows that the *complete* vertical-cue and box-cue bars were judged approximately as "gapped" as the gapped no-cue and horizontal-cue bars.

Mosteller (1951) proposed a test of the goodness of fit for this kind of scale. Conceptually, this test involves generating a set of expected proportions from the scale values and comparing them to the observed proportion data. Given the distance on the scale between a pair of stimuli, one can predict the proportion of time subjects would judge one of the stimuli as "gapped" if the two were compared. If this expected proportion differs significantly from the observed proportion, one can conclude that the scale does *not* fit the data. In other words, a failure to reject the null hypothesis is evidence that the scale accounts for the observed proportion data.

Mosteller's test was applied to the scale shown in Figure 2.12 in such a way as to increase the likelihood of rejecting the null hypothesis. That is, an attempt was made to obtain a significant statistic with the idea that if the statistic failed to be significant, there would be little room to doubt that the scale was a good fit.[1] To show how this was done, the calculations underlying Mosteller's test are described briefly.

A table of expected proportions was derived from the distances between pairs of stimuli. These proportions and the observed proportions were transformed into angles via an arcsin transformation. A χ^2 was computed according to the formula $n\sum(\theta' - \theta)^2/821$, where θ' is the angle based on the expected proportion, θ is the angle based on the observed proportion, and n is the number of observations on which the observed proportions are based. Degrees of freedom are equal to $(k-1)(k-2)/2$, where k is the number of stimuli. In the present study, $k = 8$, θ was calculated from the proportions in Table 2.1, and θ' was derived from the scale values. Thus, all of the variables involved in the calculation were fixed, except n. In the current experiment, each proportion in Table 2.1 is a mean based on 10 subjects. Each subject, in turn, made 10 judgments for each comparison. Thus, n could equal "10" or "100". As "100" increased the likelihood of obtaining a significant χ^2, it was used in the calculation. The test showed that the scale fit the data rather well, $\chi^2(21) = 16.08, p = .77$.

In summary, these results complement those of the previous two experiments. When a region of zero contrast was bounded by smoothly ending edges, that region was not perceived unambiguously as a gap. The addition of cues enhanced gap detection and caused the stimulus to appear more gap-like than when there were no cues.

Experiment 2.5

This experiment was essentially the same as Experiment 2.4. The difference was that the stimulus bars were taken from Experiment 2.3. The stimuli used were the 0% and 0.88% contrast stimuli. Inspection of Figure 2.11 shows that the gapped bar was almost always judged "gapped" more so than the complete bar. Based on the results of Experiment 2.3, the predicted ordering of the stimuli was: c0, ch, cv, g0, cb, gh, gv, and gb.

Method

Subjects. Ten University of Western Ontario introductory Psychology students took part in the study to fulfill a course requirement. Seven were

[1] In this regard, Mosteller (1951) suggested setting the significance level high ($p < .01$).

Table 2.2:

Proportion matrix for paired comparisons (Experiment 2.5)

Bar	c0	ch	cv	cb	g0	gh	gv	gb
c0	-	.53	.55	.77	.91	.87	.94	.93
ch	.47	-	.57	.59	.71	.86	.88	.94
cv	.45	.43	-	.57	.73	.87	.87	.91
cb	.23	.41	.43	-	.66	.79	.75	.94
g0	.09	.29	.27	.34	-	.67	.60	.75
gh	.13	.14	.13	.21	.33	-	.52	.73
gv	.06	.12	.13	.25	.40	.48	-	.74
gb	.07	.06	.09	.06	.25	.27	.26	-
Scale	0.0	0.2183	0.2946	0.5165	0.9987	1.2943	1.3286	1.7869

female and three were male. All had normal or corrected to normal vision. Their ages ranged from 18 to 25. None had participated in any of the previous experiments.

Stimuli. Two bars were used; they were two of the three employed in Experiment 2.3. The central region of one of the bars was of zero contrast; the other was had a 0.88% contrast edge. The four types of cues were superimposed on these stimuli, and, as in Experiment 2.4, each of these eight stimuli were paired with each of the other seven. In all other respects the design of the experiment was identical to that of Experiment 2.4.

Procedure. The procedure was the same as that used in Experiment 2.4.

Results

The mean proportion "gapped" judgments, for each comparison, calculated across all subjects are shown in Table 2.2. The table is read in the same manner as Table 2.1 of Experiment 2.4. The bottom row shows the scale values obtained with the current stimuli. The scale is graphically depicted in Figure 2.13. Mosteller's (1951) test of the goodness of fit of the scale showed that the scale did account for the data, but not as well as in the previous study, $\chi^2(21) = 29.93, p = .09$ (recall that n was chosen to increase

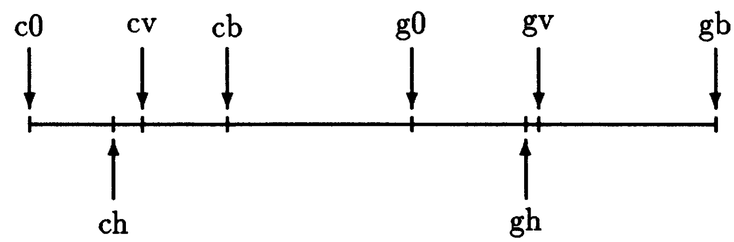

Figure 2.13: The scale for the stimuli in Experiment 2.5.

the likelihood of significance).

The scale showed that, as predicted, the no-cue stimulus was the least likely to be judged "gapped", within each set of complete and gapped bars. The appearance of the bar as a gapped figure increased as the cue changed from horizontal lines, to vertical lines, to the box. In addition, there was a clear separation between the complete and gapped stimuli in that the complete bars, as a group, were judged less "gapped" than the gapped bars. This is consistent with the results of Experiment 2.3 in which these same stimulus figures were presented individually. There, the truly gapped bar was easily seen to be so. In the present experiment, inspection of Table 2.2 shows that for comparisons between gapped and complete stimuli, subjects were fairly accurate in their judgments. Even in the case where the no-cue gapped stimulus (g0) is pitted against the most "gapped-looking" complete stimulus (cb), the gapped bar still was judged "gapped" 66% of the time. The ease with which the gap was accurately detected in this study and in Experiment 2.3 suggests, again, that it was more perceptible for this kind of stimulus.

General Discussion

The following conclusions are suggested by the data. First, it was found that steep gradients increased the detectability of a gap when there was a gap present in the stimulus. Shallow gradients, on the other hand, decreased the likelihood that a gap was perceived. These results are in accord with the idea that smoothly ending edges are good candidates for contour interpolation. Shallow gradients result in smooth endings, which do not provide strong evidence of the gap's presence. Conversely, abrupt endings did define the extent of the gap clearly and provided a definite indication of its existence.

Smooth vs. abrupt endings may be a specific instance of a general principle, namely, that unless there is some "news" in the stimulus implying that an edge ends, the visual system may decide that it does not in fact end. In these experiments, when other information cued the gap, subjects were more likely to judge that there was a gap in the bar. Thus, a smooth ending was not seen as continuing when it was marked with a variety of other edges. Furthermore, these additional edges increased the likelihood of detecting a gap even when there was none.

Reversing luminance gradients apparently made the task easier, in terms of detecting a gap, as compared to when they did not reverse. When single bars were presented, subjects judged the reversing gradient bars as "gapped" more frequently than the stimuli defined by non-reversing gradients. Also, the differences in scale value between gapped and complete stimuli were greater when the contrast gradient reversed vs. when it did not. There are a number of reasons why this may have occurred.

First, the contrast of the central region of the complete stimuli was defined differently for the two types of gradients. In the case of non-reversing gradients, increased contrast was obtained by increasing the slope of the gradient. The contrast of the central portion of the bars was therefore not a constant $x\%$, but there was only on *average* an $x\%$ difference from the background intensity. In fact, the contrast went to zero for these bars at their very centre. In comparison, the central region of the complete bars defined by reversing gradients was a *constant* $x\%$ contrast: The centre of these bars were of that contrast everywhere and never decreased to zero. Consequently, in terms of the presence of an edge through the centre, the "signal" provided by the reversing gradient stimuli may have been stronger than that of the non-reversing stimuli. This may explain the improved accuracy in Experiment 2.5 over Experiment 2.4. It was simply easier to differentiate the complete bar from the gapped in the former experiment.

Second, the greater number of stimuli in Experiment 2.2 as compared to Experiment 2.3, may have made it more difficult to remember which stimulus was which. These experiments involved presentations of single bar stimuli. Ten different gradients were used in the former experiment, and only three

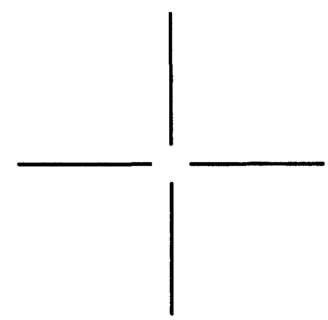

Figure 2.14: The Ehrenstein illusion.

in the latter. There is thus a greater possibility that over time, subjects would learn what the various stimuli looked like in the experiment where the contrast gradients reversed, and remember that such-and-such a stimulus was the gapped one.

Third, when the gradient reverses direction, the gap is bounded by two areas of similar contrast (in the present studies, two dark regions), and these may, via a contrast effect, enhance the brightness of the zero contrast region. These stimuli do resemble degraded Ehrenstein figures (cf. Spillmann, Fuld, & Geritts, 1976), an example of which is shown in Figure 2.14. This figure is made up of four dark lines surrounding a central gap. The centre of the Ehrenstein figure appears brighter than its background. The gap in the reversing gradient bar is also bounded by dark (on the average) regions and it has a somewhat luminous appearance. The comparable non-reversing bar does not; compare Figures 2.10 and 2.1, no-cue condition. Such an effect is a further clue to the gap's presence, which is lacking when the gradient does not reverse. Thus, the increased detectability of the gap for reversing gradients may have been the result of an induced contrast effect.

It is important to note that the interpolation effects were not all or none phenomena. A smooth ending did not lead inevitably to contour completion through the gap, it simply encouraged such an interpretation. Nor did a pair of abruptly ending vertical lines, for example, consistently induce the detection of a gap. Rather, the visual system seems to be weighing the evidence, taking into account all the information at its disposal and deciding what best accounts for the input.

This weighing of evidence may have reached a conscious level in these experiments, given the nature of the task. Subjects were asked directly whether there was a gap in a given stimulus. That kind of question probably instilled a fairly strict criterion in the subjects. When the question is put that way, it invites subjects to adopt an introspective attitude and, I expect, they could have been 100 per cent accurate if they had been given the opportunity to examine the stimuli for longer periods of time. The purpose of these experiments was not directed at how well subjects introspect, and the issue was skirted by presenting the stimuli for brief durations. In the next set of studies, the issue was avoided entirely by posing the question in an indirect manner. Subjects were no longer required to say if there was a gap in the display; rather its perception was measured by whether it functioned perceptually like a gap, or as if a contour had been interpolated across it. The purpose of this second methodology was twofold: to provide a second, different test of the hypotheses, and to investigate if interpolated contours affect visual processing in a way similar to that of real contours.

Chapter 3

Zero contrast contours and the Poggendorff Illusion

Vision is a constructive process that begins with the retinal image and, through a series of successive transformations, synthesizes a variety of percepts. This conception of visual perception is at least as old as the sensation/perception distinction. The Introspectionists, for example, proposed that *sensations* were the earliest and most primitive elements of experience from which *perceptions* were constructed, guided by the memories and judgments of the perceiver (cf. Helmholtz, 1962). The modern incarnation of this view is the "modular" or "levels" conception of perception, where each module is characterized by the function it performs, which in turn forms the basis of later processing. Thus, "low-level" vision is characterized by the extraction of basic properties of the input intensity array (eg. edges), intermediate levels by the processing of these basic properties (eg. computing the disparity among edges in a pair of images), and, at the highest level, by recognition of the objects in the scene.[1] This view of perception is certainly not without its detractors, most notably Gibson (1979), but it is not my purpose here to offer critical analysis. I assume that some version of it is correct in what follows.

A consequence of the modular approach is that one's perception at any

[1] Another distinguishing characteristic of a level is the degree to which cognitive states influence it. The Müller-Lyer Illusion is considered a low-level phenomenon because *knowing* that the two lines are of equal length does not alter one's erroneous perception of them (cf. Pylyshyn, 1984).

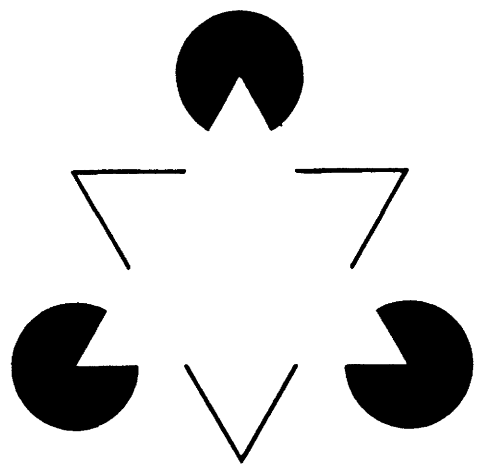

Figure 3.1: Subjective contours, from Kanizsa (1976).

given instant is relative to a level of processing. For example, we can say, without contradiction, that we *see* both a door's rectangular shape, and, at the same time, *see* that it has the form of a trapezoid. At a relatively high level, the visual system has taken account of all the cues to distance and, given the projected shape of the door, has constructed a percept of a rectangle. At a lower sensory level, the shape of the door's image is trapezoidal. What subjects "see" in an experiment will depend upon the level to which they attend. If the instructions encourage attention to relatively low level information, then subjects will attempt to report on their perception at that level. In this instance, they will report the shape of the door as a trapezoid, not as a rectangle.

The fact that information is available at one level and not at another presents a methodological problem. How can you measure something that is reported as seen or not seen depending on subjects' frame of mind? For example, the subjective contours shown in Figure 3.1 are, on one hand, quite vivid. On the other hand, it is easily seen that there is nothing in the stimulus corresponding to these contours—there is no luminance difference, no difference in texture, and no change in colour. Depending on subjects' interpretation of the question, "Are there contours here?", you could receive two different answers: "Yes, I see contours here", and "No, they are not

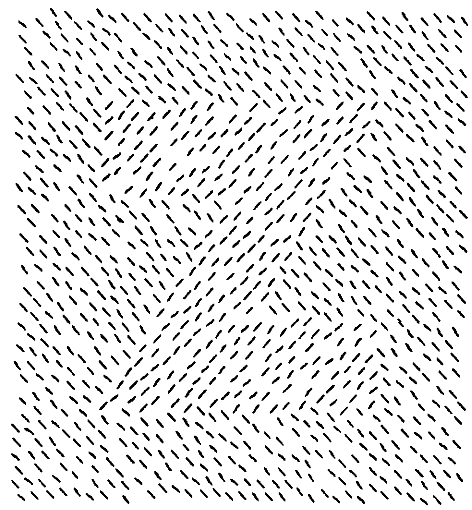

Figure 3.2: Texture contours, from Riley (1981).

really there."

One solution to this problem is to take the idea of modules seriously and assess the functional role of the contours. The rationale is if contours have been inferred at some point in the processing of the stimulus, then they must make their presence known at a subsequent stage. They are not only verbally reported by subjects, but also influence perception in some fashion. For example, Riley (1981) not only showed that texture differences elicited the perception of contours (see Figure 3.2), but went on to demonstrate that these contours were used to infer motion in apparent motion displays. He found that regions defined solely by texture contours were seen to move in a manner similar to those defined by intensity edges. In short, the texture contours were functionally equivalent, in terms of motion perception, to those defined by a luminance change.

The experiments reported in this chapter represent a functional approach to the effect of smooth vs. abrupt edge endings on contour interpolation. To the extent that smooth endings elicit completion, then even though a gap may exist in the stimulus, at some level in the visual system it is no longer perceived as such; a contour has taken its place. That contour ought to affect further perceptual processing in a manner analogous to a real contour. In essence, the experiments assessed the degree to which the gap functioned like

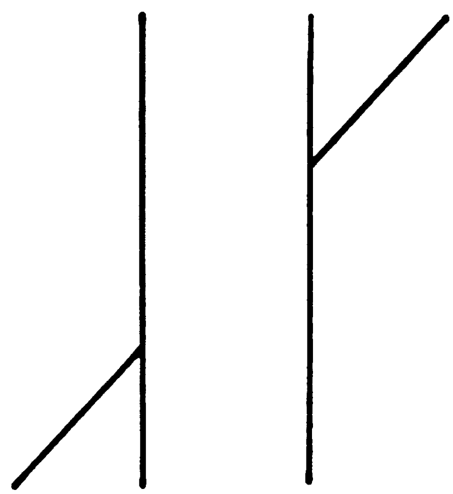

Figure 3.3: The Poggendorff Illusion.

a contour, as opposed to a region of zero contrast. The Poggendorff Illusion was employed to make this assessment. Since the reader may be unfamiliar with this illusion, and, in particular, with the relevance of contours to the illusion, a review follows.

3.1 The Poggendorff Illusion

A number of visual illusions are induced by the presence of contours; among them is the Poggendorff Illusion, illustrated in Figure 3.3. The parallel vertical contours cause observers to misperceive the alignment of the oblique transverse lines. Typically, the right transversal is perceived as too high when, in fact, the transversals fall on the same virtual oblique line, as they do in Figure 3.3. When asked to set the right transversal to a position such that it appears aligned with the left, subjects generally set it too low.

Weintraub and Krantz (1971) have varied numerous properties of the Poggendorff figure in order to determine the factors that contribute to the misperception. For present purposes, only the situation in which the parallels are vertical is considered. Under this condition, the size of the illusion depends both upon the distance between the parallels (w) and the angle made by the left transversal and the left parallel (θ). As w is increased, or as θ is made more acute, while holding the other constant, the intensity (I) of the illusion also increases. The combined effect of width and transversal angle is multiplicative. Using a least squares analysis, Weintraub and Krantz derived the following relationship among I, w, and θ:

$$I = 0.162(w/\tan\theta).$$

To what extent does the Poggendorff Illusion depend upon the presence of the parallels? Weintraub and Krantz eliminated various portions of the parallels and measured the intensity of the illusion. They found that, in general, removal of any portion of the parallels resulted in a decrement. Some portions of the parallels proved more important than others. Removal of the entire right parallel, or of that segment of the left parallel below the transversal, reduced the effect by 25 to 40 per cent. Misalignment was reduced to zero if the portion of the left parallel above the transversal had been deleted from the stimulus. In one experiment, they varied the contrast of this upper segment, and found that as the upper part of the left parallel was progressively lightened, the intensity of the Poggendorff effect decreased. When there was zero contrast between parallel and background, no effect was found.

These results suggest that misalignment of the transversals depends a great deal upon the presence of the parallels; however, this is not precisely true. For the two studies in which Weintraub and Krantz altered the parallels, the angle of the transversals was held constant; 16.7 deg in one experiment, and 24.2 deg in the other. They did not investigate other transversal angles, and their results may not have held had they done so. As a matter of fact, other transversal angles elicit an effect even when there are no parallels in the stimulus figure. Day (1973) conducted two experiments in which he varied the angle of the transversals in a "transversals-only" display. For this type of stimulus, he found that the degree of misalignment was an inverted U-shaped function of transversal angle. The effect was maximal when the transversals were oriented at 45 deg and it decreased as the transversals were rotated either towards the vertical or horizontal. In a second experiment, Day compared the effects of a standard Poggendorff display possessing parallels with one that lacked them. He held the angles of the transversals constant at 45 deg in this study. He found a significant effect with both displays, and, furthermore, the effect obtained with parallels was approximately three times that observed without parallels.

In summary, misperception of alignment does not depend solely on the existence of the parallels. It is, however, enhanced if they are present in the stimulus. Also, the relationship between the angle of the transversals and the degree of misperception is different depending on whether there are parallels in the display. With parallels, the effect is monotonically increasing as a function of the tangent of the angle; without, it is an inverted U-shaped function with a maximum at 45 deg.

3.2 Subjective Contours and the Poggendorff

If the parallels are illusory, will they induce a Poggendorff effect? Both Gregory (1972) and Kanizsa (1976) have published displays in which the upright parallels were defined by subjective contours. The transversals in their demonstrations appeared misaligned in a way similar to that found using stimuli with "objective" parallels. However, the transversals in these displays were oriented at approximately 45 deg, and given Day's (1973) results, it is not clear that the illusion observed in Kanizsa's and Gregory's figures are due to the subjective contour parallels. In order to obtain a better idea of the effect of subjective contours, it is necessary to examine them using a number of transversal angles.

There have been three attempts to systematically measure the Poggendorff effect induced by subjective contours. Only one of these was successful. First, using two transversal angles (45 deg and 51 deg), Goldstein and Weintraub (1972) measured the amount of misalignment induced by a standard Poggendorff display, a subjective contour display, and a transversals-only display. They found that the effects induced by the transversals alone failed to be significantly different from zero. The effects obtained with the subjective contours were significantly greater than those found using transversals-only, and the misperception was greater still in the case of the conventional stimulus. They concluded that subjective contours were sufficient to produce a Poggendorff effect, but that it was attenuated relative to the conventional illusion. However, contrary to the Weintraub and Krantz (1971) data, as the angle of the transversals was made more acute, the degree of misalignment decreased for the subjective contour display. This pattern, if anything,

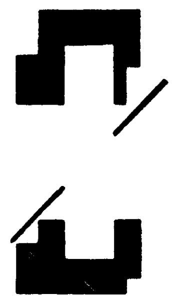

Figure 3.4: Subjective contour Poggendorff (Goldstein & Weintraub, 1972).

is similar to that reported by Day (1973) using a transversals-only display. The pattern of results obtained by Goldstein and Weintraub (1972) using the conventional display followed that previously found by Weintraub and Krantz (1971). These results suggest that subjective contours do not induce a standard Poggendorff effect.

Two comments can be made with respect to the Goldstein and Weintraub (1972) study. First, two similar transversal angles were investigated and, although their data showed an incorrect trend, it would have been more revealing had they investigated a broader range of transversal angles. Secondly, the subjective contours in their display are not very impressive (see Figure 3.4). One suspects that if the subjective contours were more vivid, then larger Poggendorff effects would be observed.

Day, Dickinson and Jory (1977) reported two studies of the effectiveness of subjective contours in terms of inducing the Poggendorff Illusion. They compared the effects obtained with a standard figure, a transversals-only figure, two subjective contour displays, one control stimulus in which the parallels were defined by dashed lines, and two other controls. The purpose of the controls was that while they had the configural properties of a Poggendorff display, they did not induce an impression of a subjective contour. Only one transversal angle was tested, that being 45 deg. They found little effect for the transversals-only display, a large effect for the conventional figure, and a substantial effect for only one of their subjective contour displays; the other induced an illusion not significantly different from that of the transversals-only. In addition, the effects obtained with the control stimuli were as great or greater than that obtained with the subjective contour figures.

Day et al (1977) suggested that a Poggendorff Illusion was not obtained with subjective contours. However, they noted that the elements used to induce the subjective contours were fixed, and that when subjects adjusted the right transversal, it was possible to set it to a position that intersected one of these elements. This was true for both the subjective contour stimulus and the control stimuli. In fact, when the transversal was placed at the position corresponding to the mean adjustments of the subjects, they found that the transversal did intersect an inducing element. Thus their experiments may have measured a standard Poggendorff, and not the effect of subjective contours. Day et al (1972) concluded that because of the intersection "between the oblique [transversal] and a semi-circular element, the role of subjective contours remains unsettled" (p. 219).

Meyer and Garges (1979) used magnitude estimation in two studies of subjective contours and the Poggendorff Illusion. The transversals were colinear in all of the displays and remained fixed while subjects estimated the size of the effect. In this way, they avoided the potential intersection of the transversals and the inducing elements. The stimulus displays included a subjective contour display, a standard Poggendorff figure, a transversals-only display and three control displays that resembled the subjective contour stimulus but did not induce their perception.

In their first experiment, the Poggendorff Illusion was explained to the subjects, and then they were asked to rate each display on a five point scale, where a higher rating denoted a greater perceived misalignment. They found a significant difference between the rankings of the subjective contour display and the transversals-only display. However, they also reported a significant effect with the control stimuli, albeit less than that induced by subjective contours.

In their second study, they did not explain the apparent misalignment induced by the Poggendorff figure the subjects. They presented the stimuli and asked the subjects to estimate the size of the effect by marking off a length on a vertical line proportional to the degree of perceived misalignment. When this methodology was used, they found significant effects only for the conventional and subjective contour displays.

Meyer and Garges' (1979) study provides some evidence that subjective contours suffice to induce the Poggendorff Illusion. Nevertheless, it would be useful to see if the Illusion can be measured using the method of adjustment. Experiment 3.1 reported below represents an attempt to do just that. The reason for undertaking this experiment was primarily to justify the methodology: If an effect could by obtained using the paradigm, then it could be used to measure the completion effects observed in the previous chapter. Past studies into the effects of subjective contours have proven ambiguous when a method of adjustment was used. If improved methods of investigation could be found, then they could be used with confidence to measure the functional role of other kinds of illusory contours.

Three variables were manipulated in Experiment 3.1. First, the angle of the transversals was varied over a broad range. The prediction was that if a stimulus induced a standard Poggendorff effect, then the amount of misalignment would continually increase as the transversals deviated from the horizontal. If the effect was that found with transversals alone, then a U-shaped function should be observed.

Second, the Poggendorff effect was studied in conjunction with three kinds of contours: objective contours, subjective contours, and virtual contours. In the objective contour condition (see Figure 3.5), the parallels were defined by an intensity edge. In the subjective contour condition, the displays were designed to elicit the perception of subjective contour parallels. In the control condition, the displays contained all of the elements used in the subjective contour condition, but were oriented in a way that precluded the perception of subjective contours. Even so, the elements were positioned in the same relative locations as in the subjective contour condition, lining up along "virtual" parallels. Note that the configuration used insured that the transversals would not intersect any contour inducing element.

A third variable, which has been found to influence the vividness of subjective contours, was included in order to determine if there is any relationship between the phenomenal impression of subjective contours and the degree to which they function like objective contours. Dumais and Bradley (1976) found, using a magnitude estimation technique, that as the retinal size of a subjective contour stimulus is reduced, the vividness of the subjective contour is increased. In the present experiment, the distance between two of the contour inducing elements was varied; the prediction was that when they were closest, the subjective contour was presumably most vivid and would cause the greatest misalignment. Furthermore, this prediction should hold only for the subjective contour condition.

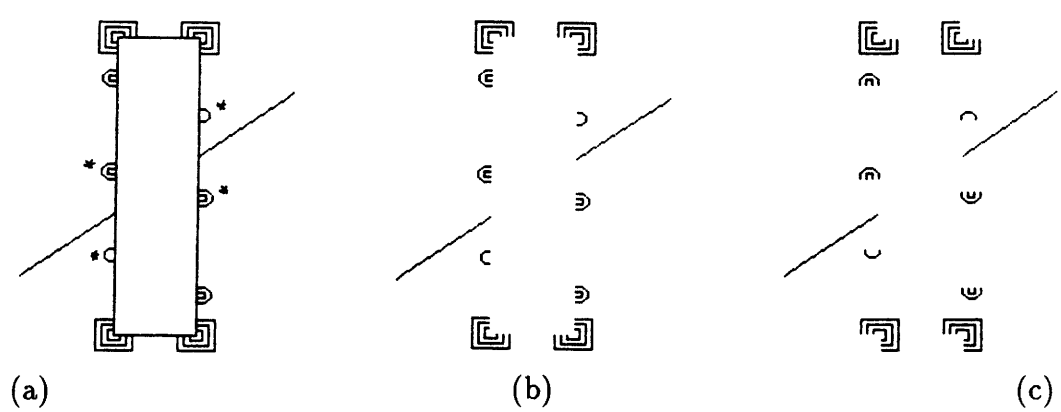

Figure 3.5: Examples of stimuli used in Experiment 3.1.

Experiment 3.1

Method

Subjects. Twelve University of Western Ontario Introductory Psychology students participated in the study for course credit. Their ages ranged from 18 to 34, and there were three males and nine females. All had normal or corrected to normal vision.

Apparatus and Stimuli. The apparatus used differed from that described in the previous chapter, and it was used only for the present study. The various displays were presented on an Apple III monochrome green monitor, model number A3M006. The viewing distance was one metre, and each pixel subtended 2.45′ of visual angle vertically by 2.6′ horizontally. The experiment was run by an Apple IIe microcomputer. A chin rest was used to steady subjects' heads.

Three kinds of contours were investigated: objective contours, subjective contours and virtual contours. Subjective contours were produced by placing contour inducing elements on the outline of a virtual rectangle (see Figure 3.5b). The dimensions of this rectangle were 7.41 deg high by 45.6′ wide. In the objective contour condition, a physical outline of the rectangle was presented in addition to the inducing elements (see Figure 3.5a). The third contour condition served as a control for the subjective contour condition. Here the displays consisted of the same inducing elements, however, they were rotated in such a manner that their alignment no longer led to the perception of subjective contours while their locations remained approximately the same as that in the subjective contour condition. An example of this type of stimulus is shown in Figure 3.5c.

The transversals were 1.73 deg in length, and were situated such that when aligned, they fell on a line that passed through the centre of the rectangle. The vertical height of the left transversal remained fixed during each trial. Subjects could adjust the vertical position of the right transversal. Five transversal angles were tested: 0 deg, 14 deg, 28 deg, 40 deg, and 51 deg; measured as the acute angle between the left transversal and a horizontal line.

The distance between certain pairs of the contour inducing elements was varied. The particular inducers involved in this manipulation were the two semi-circular shapes immediately above and below the transversals; these are marked with a "*" in Figure 3.5. Three different distances were tested, namely, 49′, 98′ and 147′. These shapes were positioned such that, for all three distances, the transversals intersected the rectangular outline (objective, subjective, or virtual) at a point midway between the "gap" defined by the shapes. This configuration was maintained throughout each trial; as subjects moved the right transversal up or down, the pair of shapes surrounding it moved an equal amount so that the transversal was always located halfway between them.

There were 15 different displays for each of the three contour condition (five angles and three interelement distances). Subjects were shown each of these displays four times resulting in a block of 60 trials for each contour condition. For two of the four presentations, at the beginning of the trial, the right transversal was situated obviously too high for alignment and on the remaining two presentations, it was situated too low. The presentation of each block of trials was counterbalanced across subjects. Within each block, the stimuli were presented in a different random sequence for each subject. A practice display was associated with each block, which consisted of the appropriate contour type, the medium sized gap, and for which the transversals were set at 40 deg from horizontal.

Procedure. When subjects arrived and were seated in front of the apparatus, the room lights were turned off in order to reduce reflection off of the screen monitor. Subjects were presented with the practice stimulus appropriate to the contour condition in which they were to be run first. They were shown which two keys on the Apple keyboard to press in order to move the right transversal up and down, and were told that their task was to move the right transversal to a position that appeared aligned with the left transversal. "Aligned" was defined by placing a straight edge on the monitor and lining up the two transversals. Before removing the straight edge, the right transversal was moved to a position of obvious misalignment. Subjects were encouraged to practise with this example stimulus in order to accustom themselves to the apparatus. They were allowed to move the transversal up and/or down as much as they liked and when they had decided that the transversals were aligned, they were told that by pressing the space bar, they would enter their response and end that trial. When subjects felt that they had had enough practice, the experiment proper began.

The first block of 60 trials was run. At the end of each trial, after the subjects had pressed the space bar, the stimulus was erased from the screen and the next display was presented. Between each block of trials, the practice stimulus appropriate to the next block of trials was shown and subjects were required to attempt this stimulus at least once. Subjects were told that the experiment was self paced and that they could take short breaks between trials and/or blocks of trials if they so desired. At the end of the experiment, they were given a short written account explaining the purpose of the study, and any questions that they had about the experiment were answered.

Results and Discussion

On each trial, the vertical position to which the subjects set the right transversal was subtracted from that corresponding to perfect alignment. If they had set it lower than perfect alignment, their deviations were designated positive, and if they had set it higher, negative. The mean of each of the four trials associated with each display was calculated for each subject. The contour condition by gap size means computed across all twelve subjects are shown in Figures 3.6, 3.7, and 3.8. The abscissa in these figures denotes the angle of the transversals, and the vertical axis represents the degree of deviation from perfect alignment. Perfect alignment is represented by zero on the y-axis. Gap size is the parameter in the figures: Each line represents one of the gap sizes.

The subjects' means were analyzed according to a three (contour condition) by five (transversal angle) by three (gap size) analysis of variance

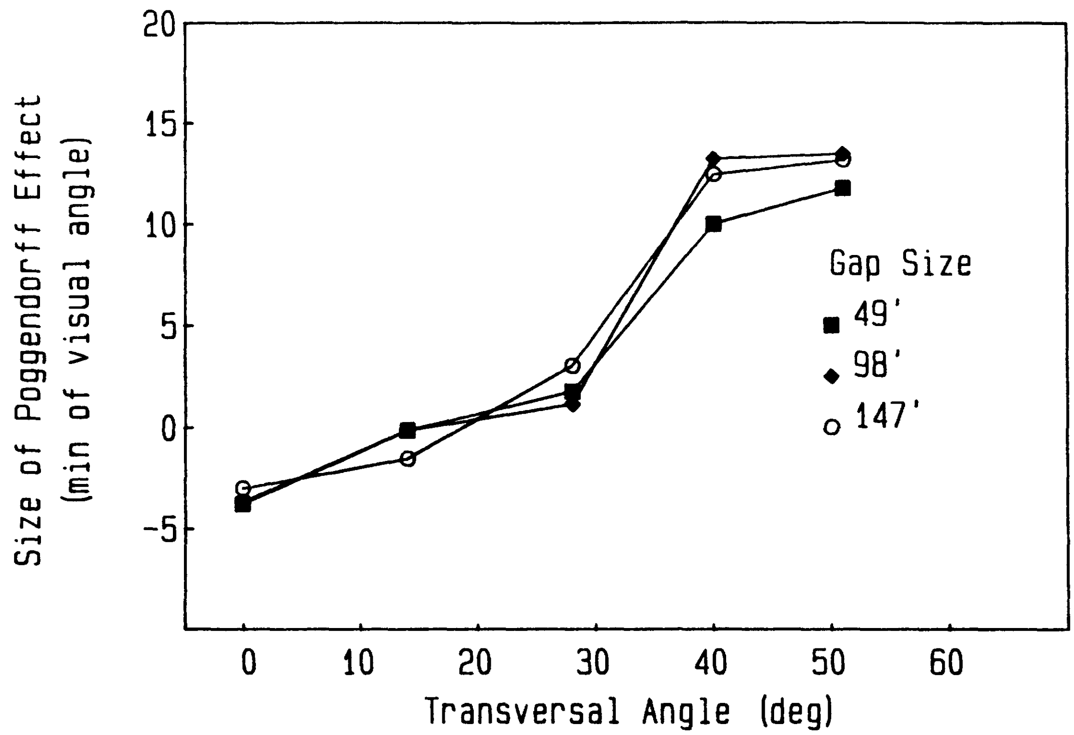

Figure 3.6: Objective contour Poggendorff effect.

with repeated measures on all factors. There was a main effect of contour type, $F(2,22) = 26, p < .001$. The objective contour condition gave rise to a greater effect (mean = 4.52′) than did the displays involving a subjective contour (mean = 1.27′), which in turn gave rise to a larger effect than the virtual contour condition (mean = −1.23′). All three means are significantly different from one another (Neumann-Keuls, $p < .05$). Inspection of Figures 3.6, 3.7, and 3.8 reveals, however, that this difference is only noticeable for the larger three transversal angles. That is, the amount and direction of subjects' deviations from perfect alignment is approximately the same for transversal angles of zero and 14 deg; whereas, for the larger angles, the means differed as a function of contour condition. This differential effect of contour type on transversal angle proved significant; there was a contour type by angle interaction, $F(8,88) = 41.58, p < .001$.

There was no main effect of gap size, $F(2,22) = 1.42$, but this factor interacted with both contour type, $F(4,44) = 2.67, p < .05$, and transversal angle, $F(8,88) = 3.21, p < .01$. The latter of these interactions suggests that gap size had little effect at small transversal angles, but had a greater influence as the transversals were set closer to vertical. The contour by gap size interaction indicates that gap size influenced subjects adjustments in only some of the contour conditions.

Examination of Figure 3.6 reveals that when subjects were presented with "standard" Poggendorff stimuli, they positioned the right transversal progressively lower as a function of transversal angle. Gap size did not affect their responses in this case. For the set of transversal angles tested, a Pearson's product moment correlation coefficient was calculated between the mean settings of the right transversal, collapsed across gap size, and the predicted settings derived from Weintraub and Krantz' (1971) formula, $I = 0.162(w/\tan\theta)$. The purpose of this calculation was to obtain some measure of the degree to which the current findings matched those of Weintraub and Krantz. The correlation was found to be 0.94, indicating a close agreement between the misalignment found using the present stimulus displays and the more conventional display used by Weintraub and Krantz. There was, however, a small difference between the present data and that of Weintraub and Krantz. It was expected that the effect would be the largest for a transversal angle of 51 deg, however, inspection of Figure 3.6 shows that the misalignment did not increase in size between 40 deg and 51 deg. Note that this occurred in all three gap size conditions suggesting that it was a real attenuation of the illusion. It is unknown why this occurred; it may have had something to do with the spatial arrangement of transversal and contour inducing elements at the largest angle.

As noted above, there was a significant interaction betweeen contour type and gap size. Misalignment was a function of gap size in the subjective contour condition. Inspection of Figure 3.7 shows that the results most closely resembled a Poggendorff effect only in the case of the smallest gap size, but, for the larger gap sizes, the function appeared similar to the one reported by Day (1973). The maximum effect was found at a transversal angle of 40 deg, for all gap sizes tested, but it did not return to zero for the larger transversal angle in the smallest gap condition. An analysis of simple simple main effects at the largest transversal angle for this contour condition proved significant, $F(2,44) = 5.97, p < .01$.

As was done in the objective contour condition, correlation coefficients were calculated between the data and Weintraub and Krantz' formula. The correlation for the means obtained in the smallest gap condition was found

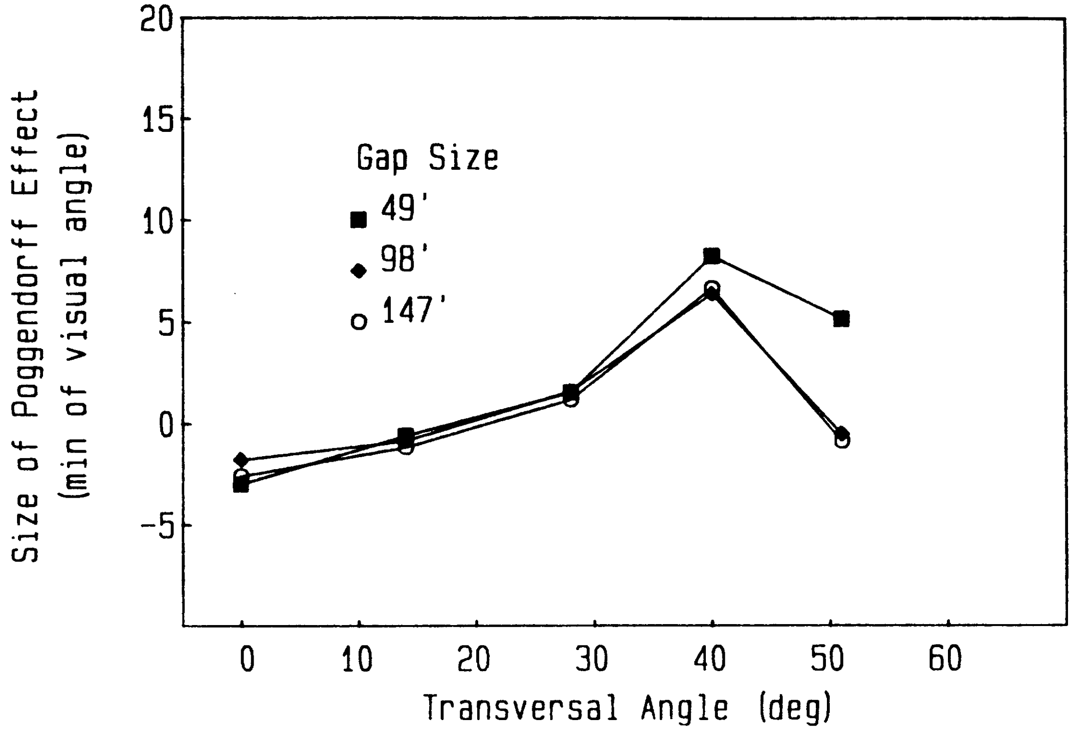

Figure 3.7: Subjective contour Poggendorff effect.

to be 0.86. For the larger gap sizes combined, the correlation was less: 0.42. These data suggest that when the impression of the subjective contours was made more vivid, by situating the inducers closer together, the illusory contours produced a stronger effect.

With respect to the virtual contour condition, little misalignment was observed for any of the transversal angles with the maximum effect occurring at an angle of 40 deg (see Figure 3.8). Gap size had no effect and the correlation between these results and Weintraub and Krantz' formula was −0.12. These results are more similar to Day's (1973) than they are to Weintraub and Krantz'.

There was one difference between the present results and Day's, which occurred in conjunction with the largest transversal angle. For some reason, subjects placed the transversal too *high* when the transversals were tilted closest to vertical, unlike Day's subjects. This finding, in conjunction with the reduced effect found in the objective contour condition at the largest transversal angle, suggests that some feature of the stimuli affected subjects'

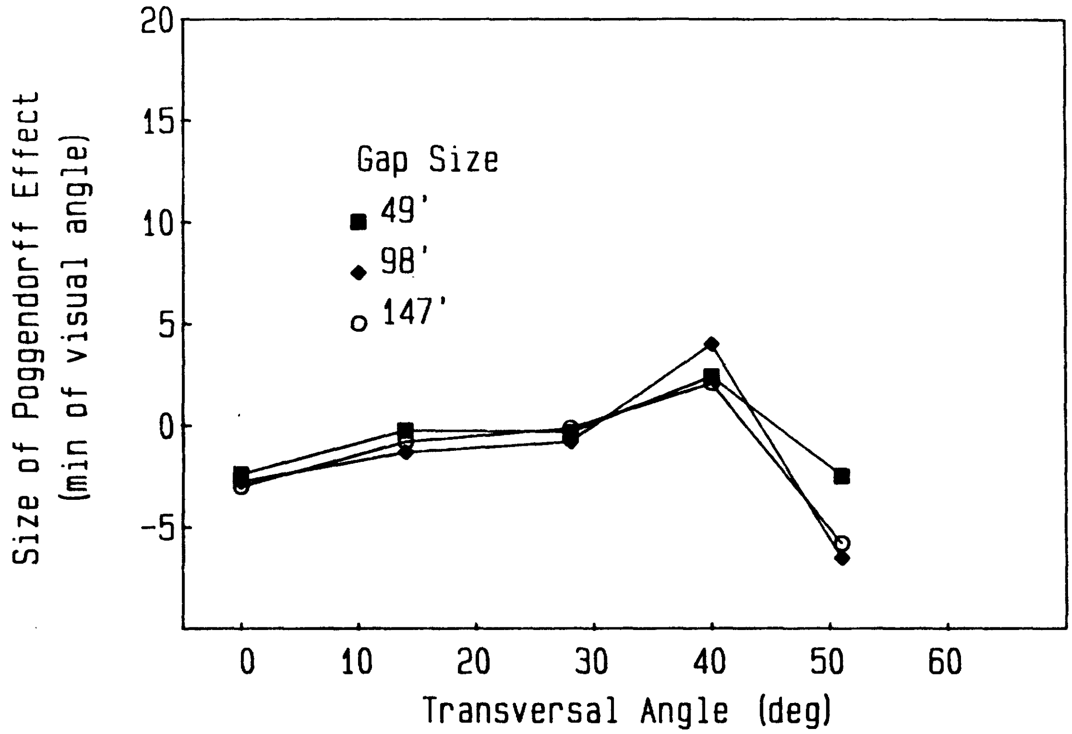

Figure 3.8: Virtual contour Poggendorff effect.

perception of alignment in a manner that contaminated the results. This unknown caused subjects to position the right transversal in a direction contrary to what is normally found using these types of stimulus figures. The implication is that the effects observed in the subjective contour condition were likewise attenuated and would have been stronger had there been no confound.

The main purpose of the experiment was to determine if the Poggendorff Illusion could be used to measure the presence of illusory contours. The results indicated that the Poggendorff effect was larger for stimuli that were likely to induce subjective contours. When the subjective contours were not that vivid (larger gap conditions), the results resembled those of Day (1973). This implies that contours were not seen, or at least not seen very well, in these conditions, and is consistent with Dumais and Bradley's (1976) finding that larger figures induce less vivid subjective contours. The gap between the inducing elements in the large gap conditions functioned like a gap and not like a contour. As the contour was made more vivid (the small gap condition), the effect was stronger.

The magnitude of the Poggendorff Illusion in the small gap subjective contour condition was not as large as that observed in the objective contour condition. This was not completely unexpected since a subjective contour is never as "vivid" as a real intensity edge (Dumais & Bradley, 1976). One would expect that even when the phenomenal impression of a subjective contour is strong, it will not have precisely the same effect on perception as a real contour. Nevertheless, this experiment argues that for the following a general rule: The greater the strength of the contour, the greater the magnitude of the Poggendorff Illusion.

3.3 Smooth endings and the Poggendorff

In the following set of experiments, the Poggendorff Illusion was used to investigate contour completion across gaps bounded by smooth vs. abrupt edge endings. Since the experimental designs and results are somewhat complex, a summary of the independent variables and main findings of each study is given immediately following.

Experiment 3.2. The effectiveness of the gap in inducing a Poggendorff Illusion was studied using gaps bounded by smooth and abrupt endings. Two types of displays were used: bars and the outlines of bars. The effects were compared to those elicited by a standard Poggendorff stimulus and a transversals-only display. Also, reversing vs. non-reversing luminance profiles were tested. The orientation of the top and bottom edges of the gap was the same as the orientation of the transversals. It was found that smooth endings induced effects greater than a transversals-only display, whereas abrupt endings did not. For smooth endings, bar stimuli induced greater effects than line stimuli. No significant differences were found between reversing and non-reversing gradients.

Experiment 3.3. There was some concern that the reduced effects obtained in the presence of abrupt endings were the result of the subjects' use of a distance estimation strategy in adjusting the transversal. That is, they could have used the clearly visible edges of the gap to guide their placement of the transversal and accurately align it with the left transversal. In Experiment 3.3, this was tested by orienting the transversals 5 deg more than the gap; this should have no effect in terms of a strategy. As a second control, framed bar stimuli were included, in which the bars were framed with black vertical edges along their entire height. For this stimulus, the strategy hypothesis predicts that the effect obtained at steep gradients would also be minimal even though there are vertical edges on either side of the gap. A third manipulation was the testing of a larger range of gradient slopes.

For non-framed bar stimuli, the effect was minimal for extremely shallow gradients, increased to a greater than transversals-only effect, and then decreased as the gradient approached a step change. At the steepest gradient, the effect remained greater than that obtained with a transversals-only display. For framed bars, the effect was maximum for the most shallow gradient, and gradually decreased as the gradient became steeper. For the line stimuli, the effect was minimal for shallow gradients, but increased as the gradient became less shallow. It did not decrease as the gradient became even more steep; the effect obtained at the steepest gradient was greater than a transversals-only display.

A problem with this study was that the type of contour (framed bar, non-framed bar, lines) was a between subjects factor, and there was an indication that some of the differences in effect were due to between the subjects. The study was replicated in Experiment 3.4.

Experiment 3.4. This study was a replication of Experiment 3.3 using a within subjects design. Three gradients were tested: none, shallow and steep. The results were essentially the same as those of Experiment 3.3. In addition, at the steep gradient, framed bars induced a greater effect than a transversals-only display, but less than a standard Poggendorff display. The last two experiments combined suggested that subjects were not using a distance estimation strategy in setting the transversal when the gradient was steep. First, the effects obtained with framed bars did not decrease to zero as the use of the strategy would predict. Second, the effects obtained for lines and non-framed bars were greater in these two studies than in Experiment 3.2. The difference between the experiments was that gap angle was equal to transversal angle in Experiment 3.2, but was not in Experiments 3.3 and 3.4. This suggests gap angle affected perceived alignment when the gap was clearly visible. Smooth endings induced effects greater than that induced by a transversals-only display in all three experiments.

Experiment 3.5. The effect of gap angle was further investigated using two different angles. One was 2 deg less than the angle of the transversals, the other was 5 deg less. Smooth and abrupt endings were crossed with framed bar, non-framed bar, and line stimuli. A transversals-only display was used as a control. The results showed that when the gap was bounded by smooth endings, the angle of the gap relative to transversal angle did not matter. When, however, the gap was bounded by abrupt endings, the effect was reduced as the angle of the gap approached that of the transversals.

Experiment 3.6. Weintraub and Krantz (1971) showed that as the contrast of the parallels was reduced, the Poggendorff effect decreased. Experiment 3.6 replicated this finding, and compared the effects induced by gapped bar and line stimuli (smooth endings) to those of low contrast Poggendorff displays. It was found that the effect induced by the zero contrast contour was approximately that induced by an intensity edge of between three and four percent contrast.

Experiment 3.2

Recall that smooth endings result from relatively shallow contrast gradients, and abrupt endings from steep gradients. If the visual system interpolates a contour when the edges end smoothly, then shallow gradients should induce a Poggendorff effect. In this case the gap should not function as a gap, but as a contour. On the other hand, steep gradients should decrease the likelihood of contour completion. Stimuli constructed from steep gradients should appear more gap-like, and should induce a relatively smaller misalignment. In this case, the gap should function like a gap and the results should resemble those obtained with a transversals-only stimulus.

Method

Subjects. Fourteen female and 10 male University of Western Ontario Introductory Psychology students participated for course credit. Their ages ranges from 17 to 33. All had normal or corrected to normal vision. None had participated in previous studies.

Apparatus and Stimuli. The equipment used in the present experiment and those following was identical to that used in the experiments reported in the previous chapter.

Two kinds of experimental stimuli and two control displays were used. The control stimuli were a standard Poggendorff display (cf. Figure 3.3), and a transversals-only display. The experimental stimulus figures were bar stimuli, similar to those used in the detection experiments, and line stimuli, consisting of only the outline of the bar stimuli. These were included for comparison: Is contour completion improved when the stimulus occupies a two dimensional region of the image (bars), as opposed to restricting it to a pairs of thin lines? The width of all of the displays was 2.25 deg of visual angle (in the transversals-only condition, the transversals were situated as if they intersected a 2.25 deg wide bar).

The contrast gradients used in the following three experiments were parabolic in shape, unlike the ramp gradients used in the detection studies. A schematic of the gradients is shown in Figure 3.9. In the present study, only two of these gradients were used; the rest were used in later studies and are diagrammed here for convenience.

Since the luminance profiles were parabolic in shape, the slope of the gradient varied as a function of position along the stimulus. The slope of the gradient was caculated as the derivative of the luminance profile at the top (and bottom) of the gap:

$$\text{slope} = dL/ds$$

where dL is the change in luminance (in cd/m^2) and ds is the change in distance along the height of the stimulus (in degrees of visual angle). Six contrast gradients are shown in Figure 3.9: 1.37, 2.72, 9.68, 13.62, 27.26, and 52.73 cd m^{-2} deg^{-1}. The luminance of the stimuli defined by the first four gradients ranged from 34–52, 24–64, 8.4–100, and 4.3–196 candelas, respectively. The luminance of the latter two ranged from 4.3–233 candelas. As can be seen in Figure 3.9, if the maximum ("white") or minimum ("black") luminance was reached, the luminance of the stimulus was set to either "white" or "black", respectively, over its remaining length. The luminance of the gap was equal to that of the background, here 43 candelas. In the present study, a relatively shallow gradient of 2.72 cd m^{-2} deg^{-1} and a steep, step change gradient were used (the latter is not shown in Figure 3.9).

Two gradient directions were tested in the experiment, one in which the luminance increased over the length of the stimulus, as is shown in Figure 3.9, and one which reversed direction and became progressively darker for the length between the bottom of the gap, and the bottom of the display. The luminance profiles of the bottom half of the reversing gradient stimuli were mirror images of their top halves. An example of a shallow, non-reversing gradient stimulus is shown in Figure 3.10, and a steep gradient in Figure 3.11.

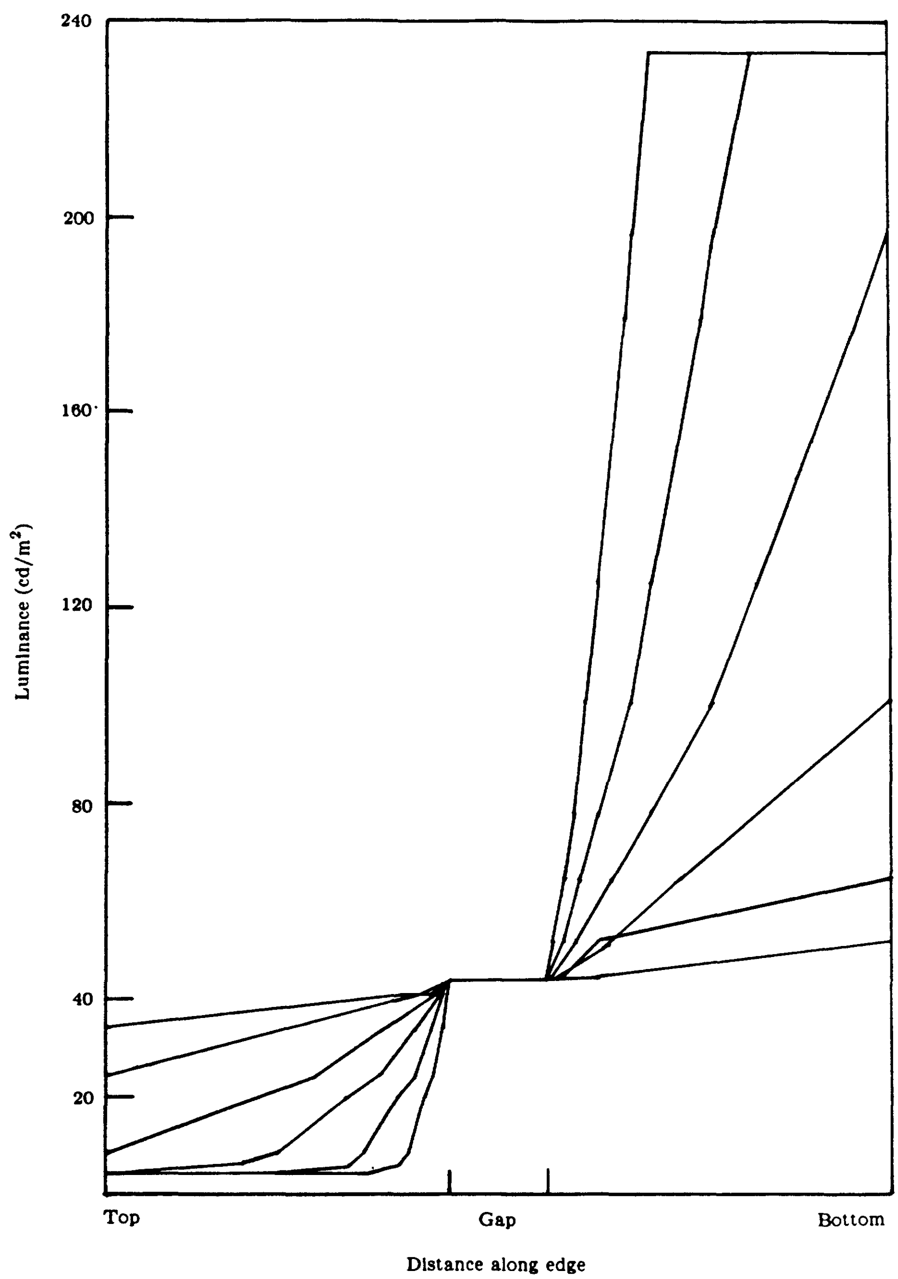

Figure 3.9: Luminance profiles of the displays used in Experiment 3.2.

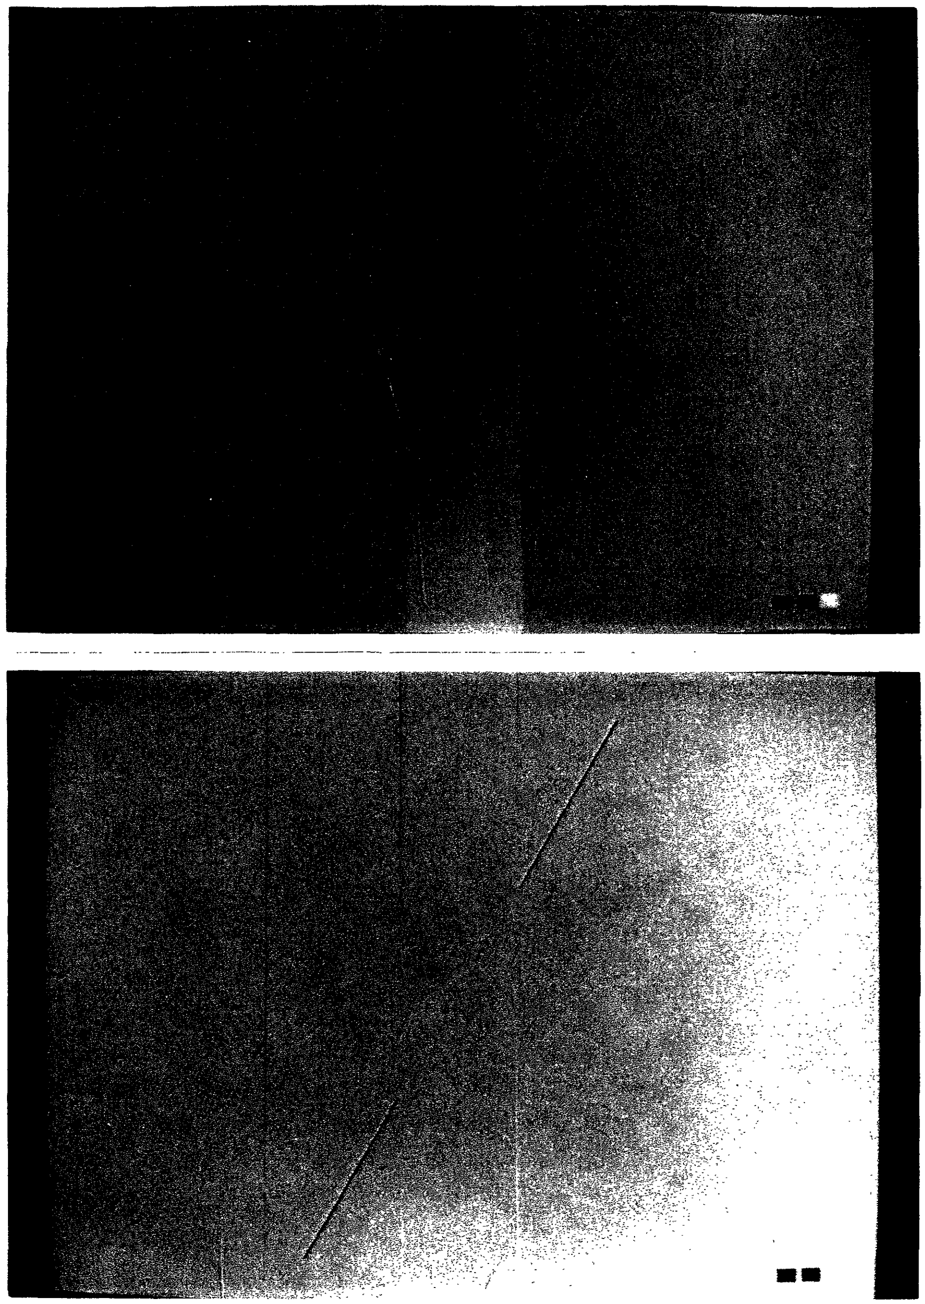

Figure 3.10: Shallow gradient stimuli (Experiment 3.2).

For all of the stimulus displays, two oblique transverse lines were displayed on either side of the rectangle. These were 3.55 deg of visual angle in length, "black" lines set 54 deg from the horizontal. Both ended at a point colinear with one of the vertical edges of the bar. The left transversal was situated such that, if continued, it would intersect the centre of the display, and it remained in that position for the duration of each trial. The right transversal could be moved up or down along the right edge of the bar and subjects were required to set this transversal to a position of apparent alignment with the left transversal.

The size of the central zero contrast region was 1.35 deg in vertical height. In order to insure that the transversals did not intersect any real edge, the gap itself was oriented parallel to the transversals. Unlike the detection experiments, it was not a horizontal cut across the centre of the stimulus, rather it was oriented at 54 deg. Assuming subjects positioned the right transversal within the gapped region, the transversal would not intersect any physically defined edge.

There were 10 different stimulus displays: two contour types (bars vs. lines) by two gradient angles by two gradient directions, plus two control stimuli. Each display was shown on six separate occasions, resulting in a total of 60 trials. For half of the six presentations, at the beginning of the trial, the right transversal was situated obviously too high for alignment, and on the remaining two presentations, it was situated too low. The sequence of presentation of the various displays was randomized for each subject. A transversals-only and a bar stimulus figure were used for practice. The luminance gradient of the practice bar did not reverse and was a step change in luminance.

Procedure. The experiment was conducted under dim lighting conditions. When subjects arrived and were seated in front of the apparatus, they were shown the first practice stimulus (transversals-only display) and were instructed that they were to move the right transversal up and/or down until it was aligned with the left transversal. "Aligned" was defined by placing a straight edge on the monitor and lining up the two transversals. Before removing the straight edge, the right transversal was moved to a position of obvious misalignment. Subjects were then shown which two keys to press on the terminal keyboard to move the line up and down, and were encouraged to practise and to accustom themselves to the apparatus. They were told that they need not worry about initially placing the line too high or too low

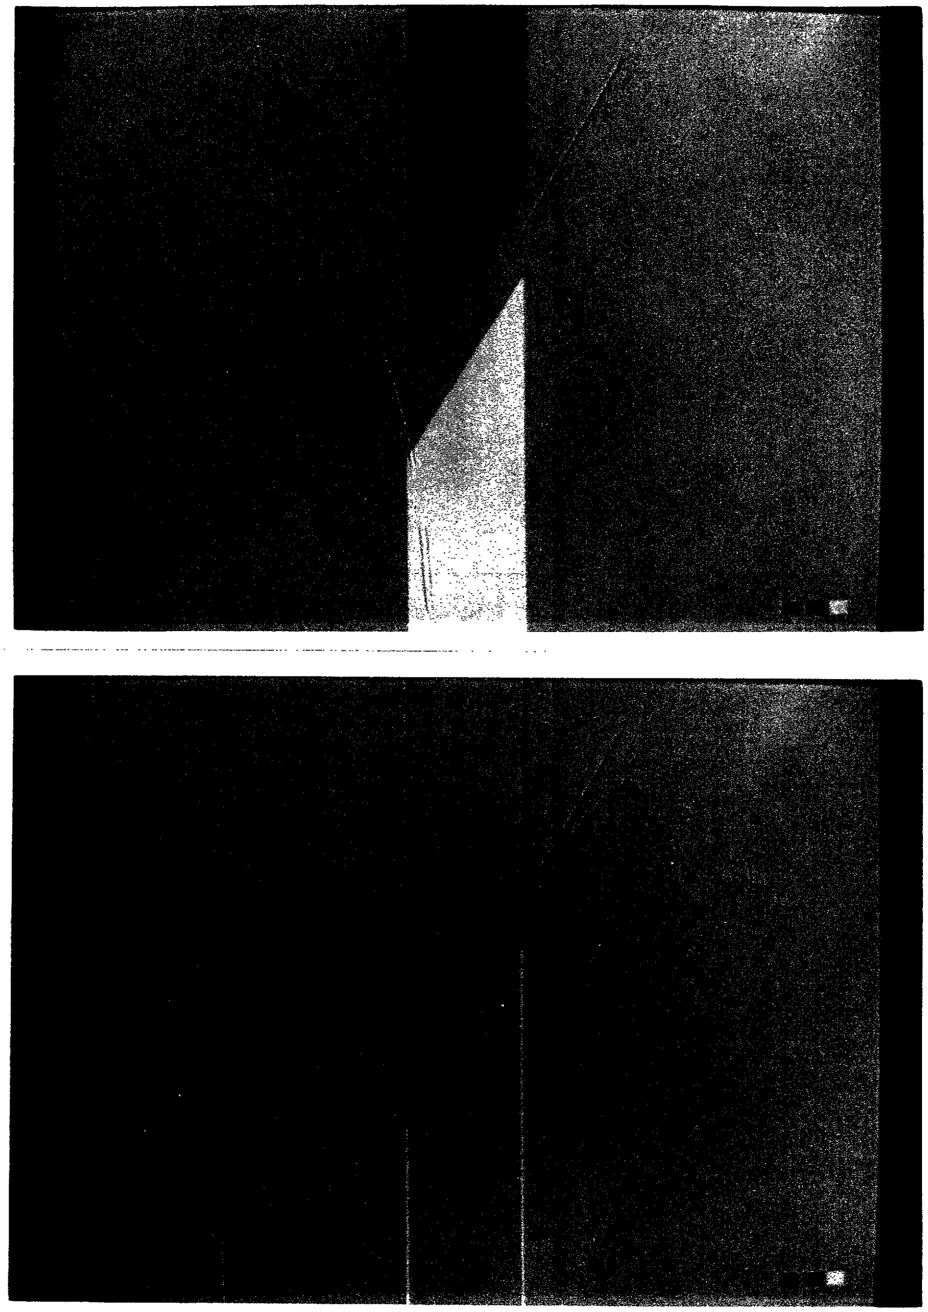

Figure 3.11: Steep gradient stimuli (Experiment 3.2).

because they could always move it back to a location they desired. When they had decided that the transversals were aligned, they were told that by pressing the space bar, they would end the trial and their response would be entered into the computer.

After practising at least once with the first practice stimulus, subjects were shown the second example stimulus in which the zero contrast region was clearly visible (step change gradient). They were then shown a number of strategies they could employ to correctly align the transversals, and they were falsely informed that none of these strategies was guaranteed to work. For example, they were told that if they attempted to visually match the distance between the top of the gap and the right end of the left transversal with a similar distance between the top of the gap and the left end of the right transversal, then that might align the two transversals, but only on some of the trials. In point of fact, this strategy would always work, provided subjects could correctly judge these distances. Subjects were misled in this manner to encourage them to focus their attention on the alignment of the transversals and not merely balance the display in some fashion. They were advised that the best way to perform the task was to look at the stimulus and ask themselves if the two oblique lines lay on the same diagonal, and if not, to move the right transversal appropriately until they did. The subjects were asked to practise with the second example stimulus at least once.

The experimental trials were then run. At the end of each trial, after the subjects had pressed the space bar, the stimulus was erased from the screen and the next display was presented. Subjects were told that the experiment was self paced, and that they could take short breaks between trials. At the end of the experiment, subjects were suitably debriefed.

Results and Discussion

Deviation from true alignment was calculated as before. The mean deviations for the bar stimuli, collapsed across subjects, are shown in Figure 3.12, and those for the line stimuli in Figure 3.13. The effects induced by the standard Poggendorff and transversals-only displays are repeated in each figure. The effect induced by the standard Poggendorff figure represents the maximum expected effect, while that produced by a transversals-only display represents the minimum. The effects obtained for these two displays were $39.63'$ and $17.18'$, respectively.

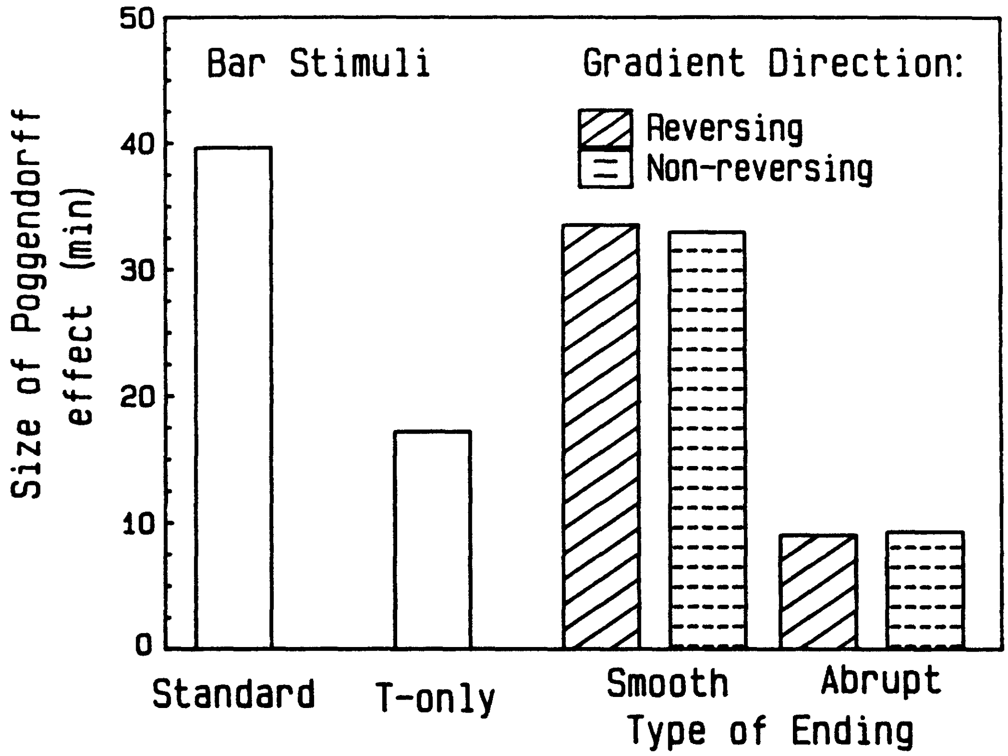

Figure 3.12: Poggendorff effect as a function of ending type and gradient direction, bar stimuli (Experiment 3.2).

Six planned orthogonal t-tests were done. They were (1) reversing vs. non-reversing gradients, (2) smooth endings vs. transversals-only, bar stimuli, (3) abrupt endings vs. transversals-only, bar stimuli, (4) smooth vs. transversals-only, line stimuli, (5) abrupt vs. transversals-only, line stimuli, and (6) smooth endings/bar stimuli vs. smooth endings/line stimuli.

Figures 3.12 and 3.13 show that there was little, if any, effect of gradient direction in any of the conditions. A t-test comparing all of the reversing gradient stimuli against all of the non-reversing stimuli failed to be significant, $t(207) = .15$. All other comparisons were made collapsing across gradient direction.

With respect to the bar stimuli, when the contrast gradient was shallow and produced smooth endings, a misalignment effect was obtained that was significantly greater than that of the transversals-only display, $t(207) =$

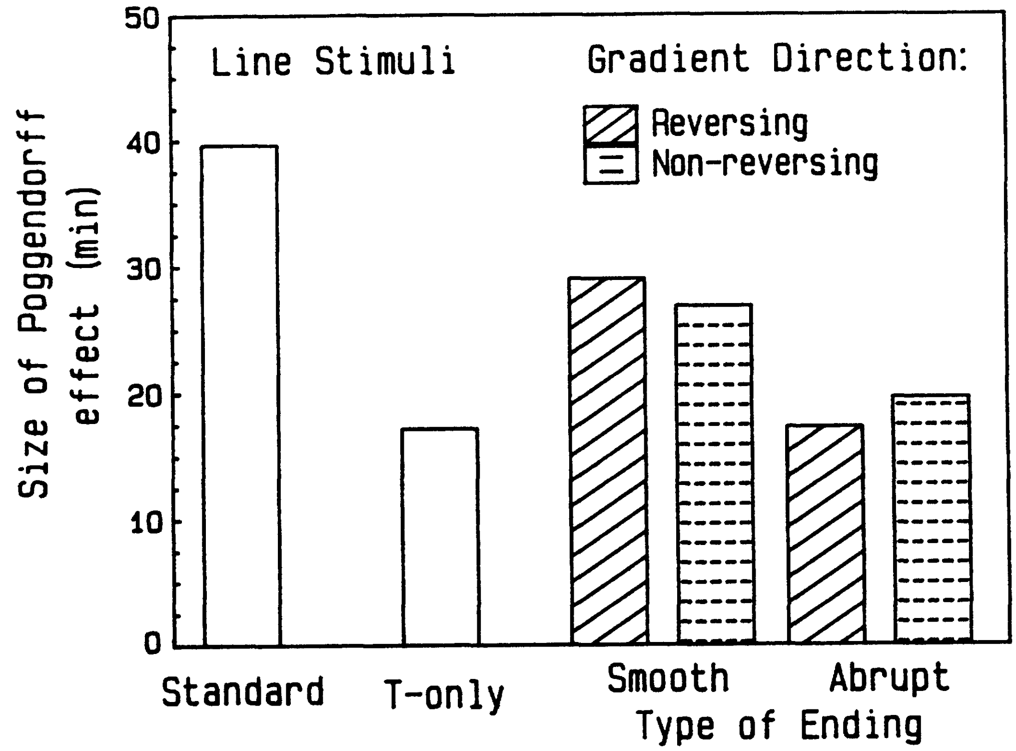

Figure 3.13: Poggendorff effect as a function of ending type and gradient direction, line stimuli (Experiment 3.2).

$8.44, p < .0001$. When the endings were abrupt, the effect was, if anything, significantly *less* than that obtained with the transversals-only display, $t(207) = -4.23, p < .0001$.

For the line stimuli, a similar pattern of results was obtained. Shallow gradients induced a larger misalignment than did the transversals-only stimulus, $t(207) = 5.64, p < .0001$, but steep gradients did not, $t(207) = .67$. A comparison of the bar vs. the line stimuli at the shallow gradients revealed that the bar stimuli induced a greater misalignment than did the line stimuli, $t(207) = 3.44, p < .001$.

The main purpose of this study was to see if gaps surrounded by smooth endings would lead to a Poggendorff effect, whereas those bounded by abrupt endings would not. The data confirmed the hypothesis. Effects obtained with smooth gradients were larger than those observed with a transversals-only display. Steep gradient displays induced effects that were either equal to or less than a transversals-only effect. The finding that bars at a steep gradient actually induced less misalignment than the expected minimum is due perhaps to the better definition of the gap in these stimuli. That is, the full extent of the gap was clearly evident in these stimuli as compared to the line stimuli, where only the endpoints of the gap were visible. These results imply that the "clarity" of the gap influenced subjects positioning of the transversal. It is possible that when the gap was clearly defined, subjects used a strategy, contrary to the instructions. However, it is not clear why they did not use it equally well with both types of stimuli (lines and bars). The effects of gap clarity on alignment were studied in greater detail in the experiments that follow.

No effect of gradient direction was found in the present study, unlike the findings of the detection studies. The reason for this difference is discussed in the general discussion at the end of the chapter.

Bars were better stimuli than lines in terms of inducing a Poggendorff effect when the gradient was shallow. This does imply that completion is better indicated by relatively large regions that end smoothly as opposed to smaller regions. Still, the lines were enough to induce a greater effect than a transversals-only display.

Experiment 3.3

The previous experiment provided evidence that smooth gradients led to contour completion. When subjective contours were used to induce a Poggendorff Illusion (Experiment 3.1), it was found that the more vivid the illusory contour, the greater the misalignment. These two results together suggest that as the gradient is made more steep, the visual system is less likely to fill in the gap; hence, the Poggendorff effect should decrease in size as the slope of the gradient is increased. The previous experiment tested only two gradient slopes. The experiment following sampled a broader range.

There was evidence in the previous experiment that something more than a simple transversals-only effect is observed when the gap was clearly visible, and a possible explanation is that subjects used a distance estimation strategy. That is, subjects could have set the right transversal such that the distance between it and the top (bottom) of the gap equalled that between the left transversal and the bottom (top) of the gap. If they had done that, then there should have been little misalignment in the steep gradient condition, as was found.

The strategy hypothesis was evaluated in the following experiment in two ways. First, the stimuli were constructed such that the orientation of the gap was no longer equal to the tilt of the transversals. If subjects were using a strategy, the results should be the same as in Experiment 3.2, since the same strategy can be used. Gap orientation is irrelevant to a distance estimation strategy. If, on the other hand, the tilt of the gap affects subjects' *perception* of alignment, the results of the current experiment should be different from the previous. Second, another type of display was included in which the bars were "framed" with a pair of vertical lines. These latter stimuli can be thought of as a combination of a standard Poggendorff stimulus and the bar stimuli used in the previous experiment. If subjects estimate distances among components of the display when the gap is clearly visible, then even though there are real contours intersecting the transversals, no Poggendorff effect should be observed. If subjects do not use a strategy, then significant misalignment should result.

Method

Subjects. Thirty-six introductory Psychology students took part in the experiment for course credit. Twenty-five were female and 11 were male; their ages raged from 18 to 22. All had normal or corrected to normal vision. None had participated in any previous experiments. The 36 subjects were randomly assigned to three groups of 12, each of which was shown one of the three sets of stimuli described below.

Stimuli. Three types of stimuli were used: framed bars, non-framed bars and lines. The width of the stimuli was 1.97 deg of visual angle. An example of a steep gradient, framed bar stimulus is shown in Figure 3.14. Eight gradient slopes were tested: 0, 1.37, 2.72, 9.68, 13.62, 27.26, 52.73 cd $m^{-2}\,deg^{-1}$, and a step change profile. These are diagrammed in Figure 3.9. Note that for both the non-framed bar and line stimuli, a 0 cd $m^{-2}\,deg^{-1}$ gradient is equivalent to a transversals-only display. In the case of framed bars, this gradient slope results in a standard Poggendorff stimulus. In all cases, the zero contrast gap in these displays was tilted 60 deg counterclockwise from the horizontal. Its size was 1.35 deg of visual angle, as before. The transversals were oriented at 65 deg from the horizontal.

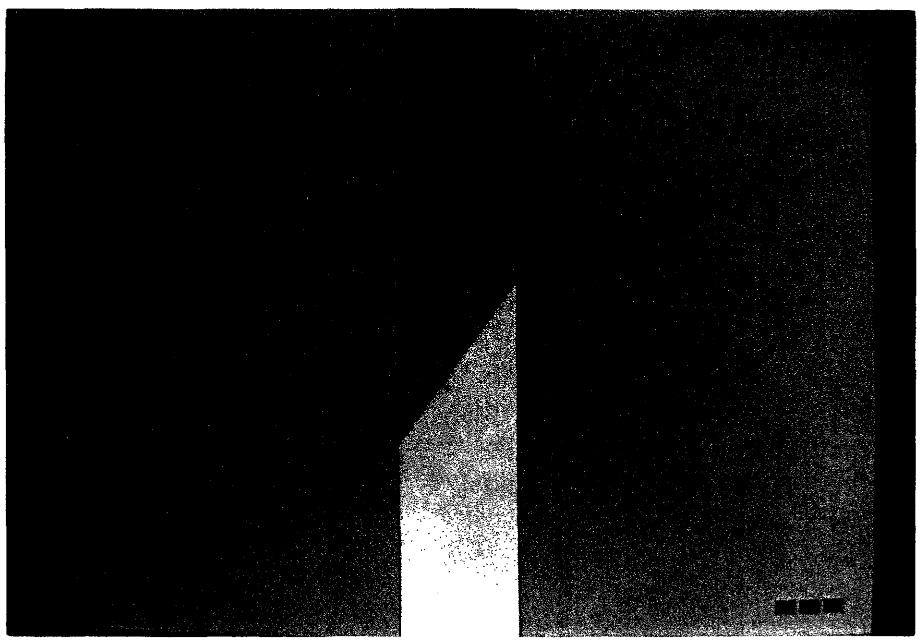

Figure 3.14: Steep gradient framed bar (Experiment 3.3).

Each group of subjects were shown only one of three sets of stimuli; either the framed bars, the non-framed bars, or the lines. Each of the eight displays in the set was presented on eight separate occasions resulting in a total of 64 trials for the entire experiment. The stimuli with luminance gradients of 9.68 cd m^{-2} deg^{-1} and the step change gradient were used for practice.

Procedure. The procedure was similar to that used in the previous experiment. The main difference was that each group of subjects viewed only one of the three types of stimuli.

Results

Deviation from perfect alignment was calculated as before. The mean gradient by contour deviation scores, computed across all subjects, are shown in Figure 3.15. The x-axis indicates the slope of the luminance gradient and the y-axis shows the size of the Poggendorff effect. The top curve is a plot of the data obtained using framed bars, the middle curve shows the results obtained with the line stimuli, and the bottom curve shows the mean misalignment scores for the non-framed bars condition.

The means were analyzed according to a three (contour type) by eight (gradient) analysis of variance, with repeated measures on the latter factor. There was a main effect of contour type, $F(2,33) = 7.31, p < .01$, of gradient, $F(7,231) = 8.92, p < .001$, and the interaction between contour type and gradient was also significant, $F(14,231) = 13.88, p < .001$. It should be noted that there was an unexpected difference between the 0 cd m^{-2} deg^{-1} non-framed bar and line stimulus conditions (see Figure 3.15, "0" on the x-axis). These two conditions represent the same stimulus figure, a transversals-only display, and should have induced the same degree of misalignment. However, the difference between them proved to be significant by a Neumann-Keuls procedure ($p < .01$) and implies that some of the difference between the three stimulus conditions is a function of the between groups design of the experiment. For this reason, no further between group comparisons were made; the remaining comparisons reported below were were within a contour condition.

Consider first the data in the framed bar condition. The stimulus possessing no contrast gradient in this contour condition was a standard Poggendorff display and it was expected that the results obtained here would be maximal. A large effect was observed for this stimulus. The effect decreased as

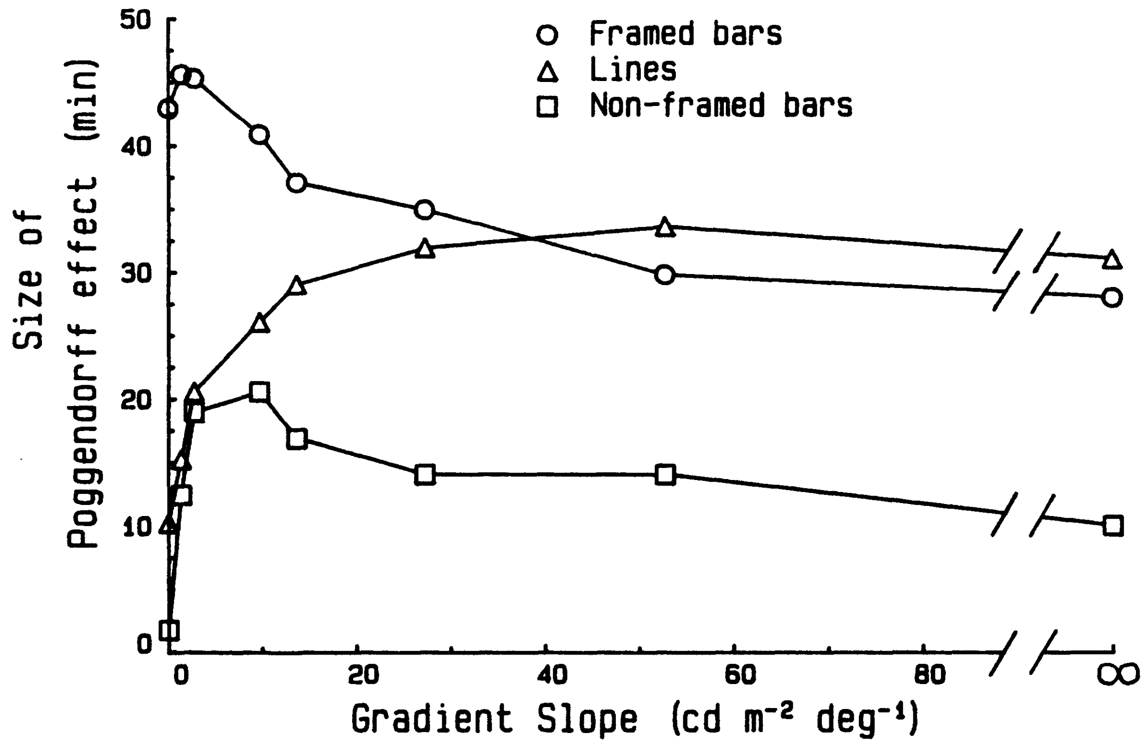

Figure 3.15: Poggendorff effect as a function of gradient slope and contour type (Experiment 3.3).

the gradient became steeper (see Figure 3.15). A Dunn's test comparing the shallowest gradient to the steepest showed that the 42.9′ effect associated with the conventional Poggendorff stimulus was significantly greater than the 27.93′ of misalignment observed in conjunction with the framed rectangle having the steepest gradient ($t'D(5, 231) = 5.18, p < .01$).[2] Unfortunately, it was unclear whether the effect obtained with the steepest gradient was

[2] All told, six comparisons were made: the one described as well as three comparisons involving the 0, 2.72 cd $m^{-2}\,deg^{-1}$, and step change gradients in the non-framed bars condition; and both of the step change and 2.72 cd $m^{-2}\,deg^{-1}$ gradients compared to the 0 cd $m^{-2}\,deg^{-1}$ gradient in the lines condition.

greater or equal to that obtained with a transversals-only display, since that is a between groups comparison.

For the non-framed bars, the Poggendorff effect was relatively small when there was no contrast in the display. The effect rapidly increased for shallower gradients, and was gradually attenuated as gradient slope increased. It did not completely vanish at the steepest gradient. In order to determine if the Poggendorff effect was significantly different from that obtained from a transversal-only display, Dunn's procedure was used comparing the means in the 0, 2.72 cd m^{-2} deg^{-1}, and step change gradient conditions. The effect observed when the gradient was shallow (2.72 cd m^{-2} deg^{-1}) was significantly greater than both that observed with no gradient, $t'D(5,231) = 5.99, p < .01$, and that involving a step change in luminance, $t'D(5,231) = 3.09, p < .01$. The effect was also greater when the gradient was steepest as compared to the transversals-only display, $t'D(5,231) = 2.86, p < .01$. Unlike the previous study, the amount of misalignment was significantly greater for the abrupt endings condition as compared to a transversals-only display for the bar stimuli.

The responses obtained in conjunction with the lines-only stimuli were different from the effects obtained with either of the bar conditions described above. The degree of misalignment induced by these stimuli was relatively small for the transversals-only stimulus, rapidly increased for shallow luminance gradients, and then approached an asymptote as the slope of the luminance gradient increased. The $31'$ effect found in conjunction with the steepest gradient stimulus was significantly greater than the $10.21'$ effect produced by the transversals-only display, $t'D(5,231) = 7.2, p < .01$. Also, the effect produced by the 2.72 cd m^{-2} deg^{-1} gradient was greater than that found using the transversals-only display, $t'D(6,231) = 3.58, p < .01$.

An annoying aspect of these data is that the individual differences in subjects' responses were so large as to disallow between group comparisons. In order to compare the effects induced among the three types of stimuli, another experiment was run using a *within* group design. This experiment is reported immediately following. The present results are then considered in conjunction with those of Experiment 3.4.

Experiment 3.4

Method

Subjects. Twelve University of Western Ontario first year Psychology undergraduate students took part in the study for course credit. None had participated in previous experiments. Six were female and six were male. Their ages ranged from 18 to 24. All had normal or corrected to normal vision.

Stimuli. The entire range of gradient angles used in the previous experiment was not tested here. Rather, a subset of three of the gradients was used. These were the 0, 2.72 cd $m^{-2}\,deg^{-1}$, and step change gradients. As before, three types of contour were investigated, namely, framed bars, non-framed bars, and lines. This resulted in eight stimulus displays: a transversals-only display, a standard Poggendorff display, and six displays involving contrast gradients of different slopes.

Procedure. All subjects were presented with each of the eight stimulus displays. In all other respects, the procedure was identical to that employed in the previous experiment.

Results and Discussion

The mean responses are shown in Figure 3.16 for each of the eight displays. The results were similar to those obtained in the previous experiment. For the non-framed bars, the effect induced by a 2.72 cd $m^{-2}\,deg^{-1}$ gradient was significantly greater than that obtained with a transversals-only display, $t(77) = 8.11, p < .001$. The effect was also greater than that obtained with the step change gradient, $t(77) = 3.43, p < .01$. Similarly, the effect obtained with the line stimuli at a gradient slope of 2.72 cd $m^{-2}\,deg^{-1}$ was significantly greater than the one observed with the transversals-only stimulus, $t(77) = 3.99, p < .001$, but less than the effect obtained with a step change gradient, $t(77) = -2.97, p < .01$. For all three contour conditions, including the framed bar condition, the effects observed at the steepest gradient were significantly greater than that obtained with the transversals-only stimulus (all t's significant at $p < .001$).

Comparing the results among the three contour types, at a gradient slope of 2.72 cd $m^{-2}\,deg^{-1}$, the mean Poggendorff effect was significantly larger for the framed bar as compared to the non-framed bar, $t(77) = 4.93, p < .001$.

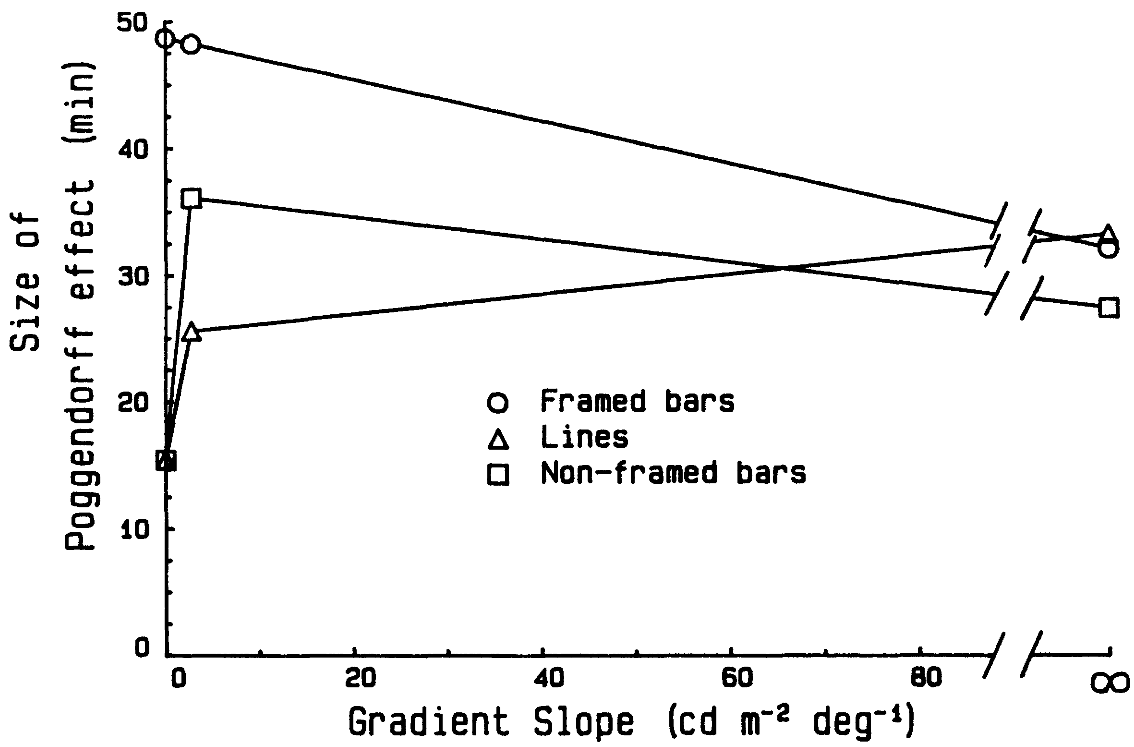

Figure 3.16: Poggendorff effect as a function of gradient slope and contour type (Experiment 3.4).

The effect for the non-framed bar was greater than that for the lines stimulus, $t(77) = 4.12, p < .001$. Thus, for shallow gradients, the pattern of results was the same as that obtained in Experiment 3.2. Bar displays induced a greater effect than lines alone. A framed bar induced the largest effect.

The purpose of the last pair of experiments was twofold: First, to determine if the misalignment of the transversals decreased as the slope of the contrast gradient increased, and, second, to test the hypothesis that subjects used a strategy when the top and bottom of the gap was clearly visible. With regard to the former, the size of the Poggendorff effect did decrease as gradient slope increased when bar stimuli were used. In contrast, this did not occur for the lines stimuli—the effect at the steepest gradient was the largest obtained in this contour condition. This result is inconsistent with that observed in Experiment 3.2 where the line stimuli induced, essentially, a transversals-only effect. Nevertheless, recall that there was a difference between the two experiments, namely, the orientation of the gap equalled to the angle of the transversal in Experiment 3.2, but was 5 deg less in the last two experiments. This suggests that gap angle is a contributing factor to the perception of alignment when the gap is clearly visible.

The data did not support the hypothesis that subjects employed a distance estimation strategy when the gap was clearly visible. In the first place, if they had, then the effects would have equalled those obtained in the transversals-only condition even when there were physical edges framing the gap. It is true that the effects were attenuated in the framed bars condition, but they remained greater at the steepest gradient than a transversals-only effect. It appears that subjects' settings were influenced both by the presence of the frame and the perception of the gap. Furthermore, for non-framed bar and line stimuli, when the orientation of the gap equalled transversal angle (Experiment 3.2), the effects obtained were less than a transversals-only effect. When there was disparity between gap angle and transversal angle (Experiments 3.3 and 3.4), the effects were greater than a transversals-only effect. The results obtained with the framed bars, non-framed bars, and line stimuli at the steep contrast gradient all imply that subjects were not using a strategy. If a strategy were used, the effect should have decreased always to that obtained with a transversals-only display (or less). That did not occur; instead the effects were a function of the difference between gap and transversal angle. The data suggest that the *shape* of the gap influenced subjects' *perception* of alignment at the steep gradient. At shallow gradients this shape is poorly defined since the contrast of the edges forming the top and bottom of the zero contrast region is extremely low. As the slope of the gradient increases, these edges increase in clarity. One hypothesis is, therefore, because it was not seen at shallow gradients, gap shape did not influence subjects' perception of alignment. For steeper gradients, since it was clearly visible, the shape of the gap may have influenced subject's responses. This hypothesis was tested in Experiment 3.5.

Experiment 3.5

Method

Subjects. Two subjects participated in this experiment, the author (JS) and one other (RS). Their ages were 28 and 29, and both had normal or corrected to normal vision. Subject RS knew of the Poggendorff Illusion but was unaware of the specific purpose of the experiment.

Stimuli. Two gradients were used: a step change in luminance, and a relatively shallow one (2.72 cd m^{-2} deg^{-1}). In addition, two gap angles were tested: 60 deg and 63 deg tilted counterclockwise from the horizontal. These angles are very similar to one another, although the difference between them is detectable if they are placed side by side. In Figure 3.17, the 63 deg gap angle is shown on the left, and the 60 deg on the right. The angle of the transversals, 65 deg, differed from both of the gap angles. Framed bars, non-framed bars, and lines stimuli were tested. A transversals-only display served as a control.

Procedure. The method of constants was used to obtain alignment thresholds. On the first occasion that any stimulus figure was tested, an estimate of the threshold was obtained from the subjects using the method of adjustment. Two adjustments were made, one in which the right transversal was initially too high and a second in which it was initially too low. The mean of these two adjustments served as the central value of a set of seven constant right transversal positions. Three of these were above and three were below the estimated point of alignment. The angular distance between the seven positions was $4.06'$ of visual angle.

There are two reasons why the constant stimuli were based on an adjustment threshold. First, for the Poggendorff Illusion the difference between perceived and physical alignment is not near zero, as it is for other thresholds. For example, if one desired to measure the difference threshold for weight relative to a standard weight of 45g, one would centre the constant weights about 45g. It is likely that the weight *jnd* includes 45g within its bounds in this instance. For the Poggendorff Illusion, it is very unlikely that the alignment *jnd* is anywhere near true alignment. When determining a difference threshold, it is desirable to confine the measurements to stimuli that are near threshold, since stimuli that are obviously misaligned are not informative in terms of what the threshold is.

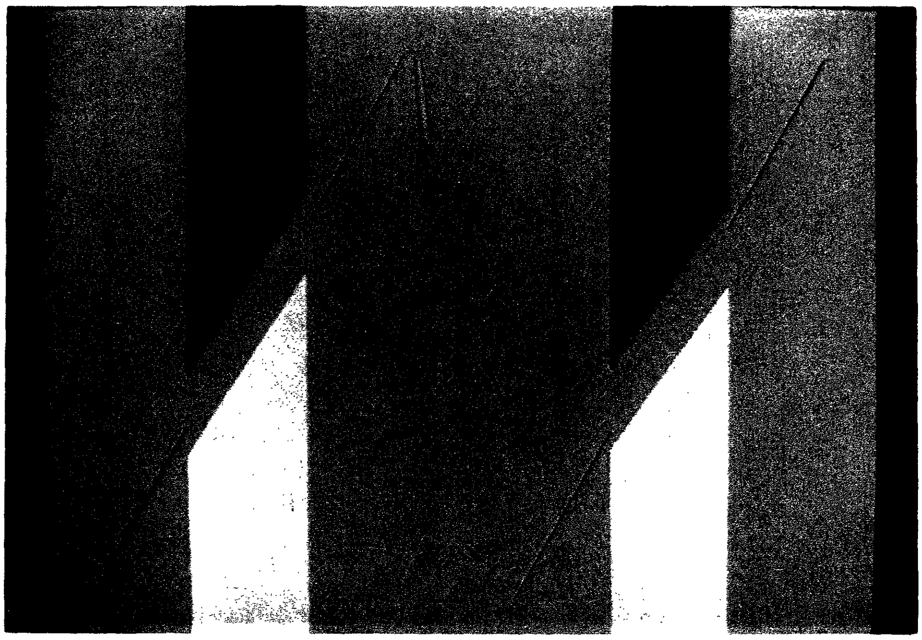

Figure 3.17: Gap angles used in Experiment 3.5.

Second, there are individual subject differences with respect to perceived alignment (cf. Experiment 3.3). What one subject sees as aligned may or not be what another considers aligned. For these two reasons, both subjects in the present experiment were required to produce an estimate of perceived alignment on which to base the constant stimuli.

Subjects were run in blocks of 140 trials on separate days. During any particular session, a pair of stimuli was tested, each of which were of the same contour type and possessed the same gradient characteristics, but differed from one another in terms of their gap angles. The pair of stimuli and the seven transversal positions were shown randomly intermixed. On a given trial, the subjects were required to view the stimulus and judge whether the right transversal was higher than the point of perceived alignment, pressing the "y" key on the keyboard if they judged it above, and the "n" key otherwise. Each of the seven transversal positions were viewed on 10 occasions per session. Each stimulus pair was tested in two separate sessions; consequently, each transversal position was tested 20 times in total.

Results

The observed proportions of "above" judgments for each transversal position were converted to z-scores and a linear regression was calculated between these scores and the transversal positions. The regression equations were used to estimate the vertical height at which the subjects judged the transversal above alignment 50% of the time, and this point was stipulated the threshold of alignment. The *jnd*'s were one standard deviation above and below the 50% point.

The thresholds and *jnd*'s obtained from the two observers are shown in Figures 3.18 and 3.19. The x-axis of these figures designates the type of contour tested; the y-axis represents the size of the Poggendorff effect. The filled symbols (the two upper curves) show the alignment thresholds obtained with the shallow contrast gradient, for all three contour conditions. Specifically, the filled triangles denote the effects obtained with a gap angle of 60 deg; and the filled squares, the 63 deg gap. Open symbols (the two lower curves) depict the misalignment effects obtained using the steep gradient. Again, triangles are associated with the 60 deg gap angle, and squares with the 63 deg gap angle. The transversals-only stimulus is denoted in the figures by "Control".

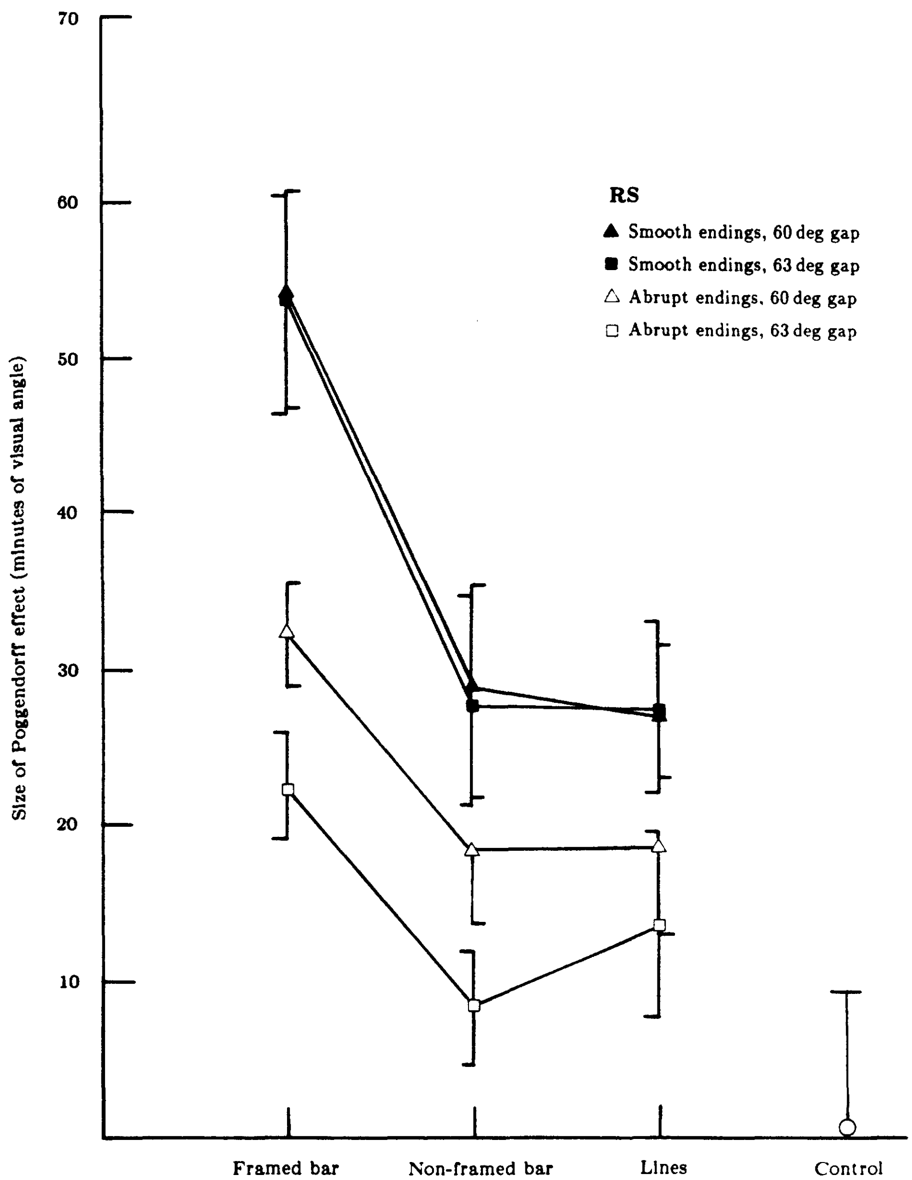

Figure 3.18: Alignment thresholds and *jnd*'s for subject RS (Experiment 3.5).

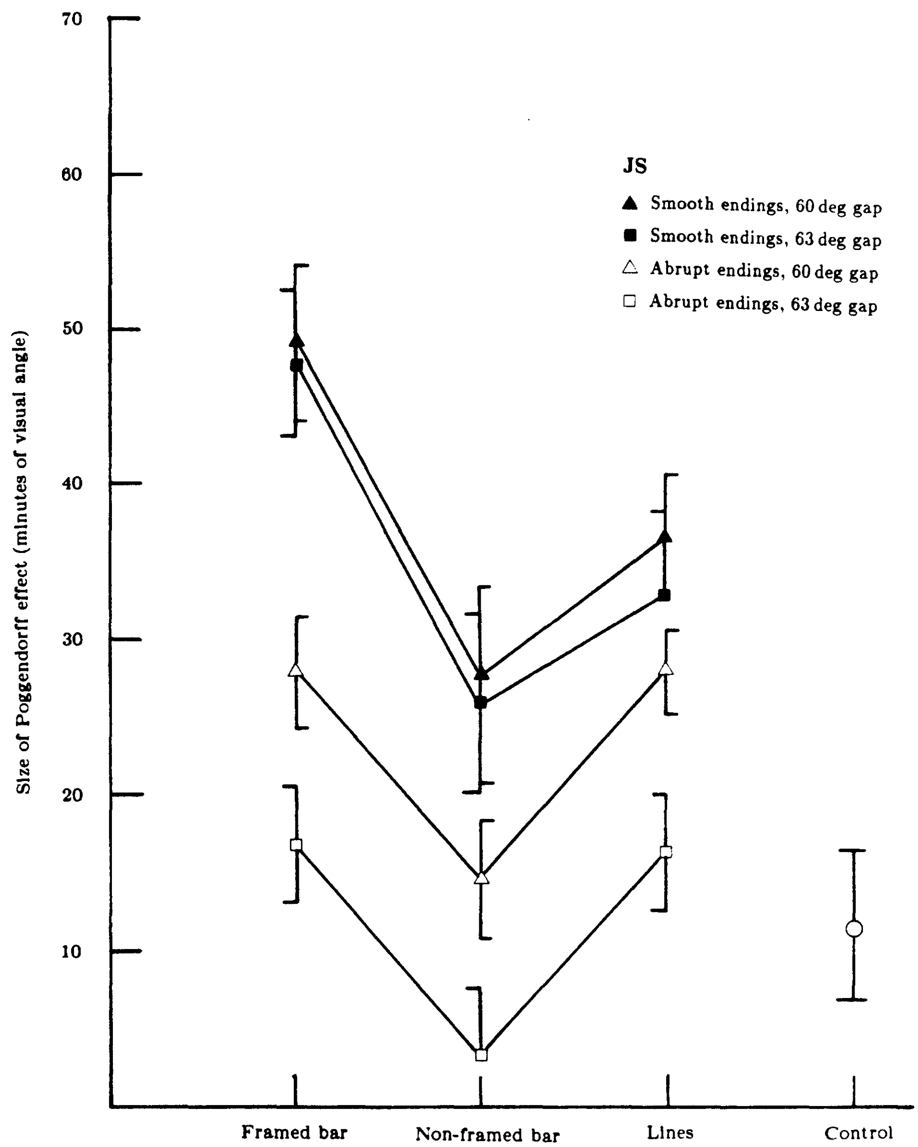

Figure 3.19: Alignment thresholds and *jnd*'s for subject JS (Experiment 3.5).

It is clear that alignment threshold was independent of gap angle at the shallow gradient: The two upper curves are virtually superimposed. For the steep gradients, however, the thresholds for the two gap shapes were quite different—perceived alignment depended a great deal upon the tilt of the gap. Comparing the open triangles with the open squares, the Poggendorff effect differed, in the case of RS, an average of $8'$ of visual angle, and in the case of JS, an average of $10'$, for a difference of only 3 deg in gap angle. The difference in thresholds depended upon the angular disparity between the gap and the transversals. The amount of misalignment tended towards that of the transversals-only display as the gap angle approached that of the transversals. Note also that gap shape influenced perceived alignment, at the steep gradient, even when there were physical edges framing the gap. It seems, therefore, that when the gap is clearly visible, its shape is a major determinant of subjects' judgments of alignment.

The effects observed in the present experiment are consistent with the data of the previous three experiments. When the parallels were defined by smoothly ending edges, a misalignment of the transversals was obtained, which was greater than that observed with a transversals-only figure. When the parallels ended suddenly, the zero contrast region was made more definite, and the amount of misalignment decreased. This decrement in misalignment was obtained even when the zero contrast region was framed by physical contours. In addition, the size of the effect depended upon the angle of the transversals relative to the angle of the gap, but only when the parallels ended abruptly above and below the gap; no such dependency on gap shape was found when the contrast gradient was shallow.

Experiment 3.6

It was noted earlier that Weintraub and Krantz (1971) discovered a relationship between contour contrast and the Poggendorff effect. They found that the size of the Poggendorff effect was directly proportional to the contrast of the left vertical parallel above the transversal: The Poggendorff effect decreased as the parallel was made less distinct. When contrast was zero, the effect was zero.

A consistent finding in the experiments reported in this chapter is that the Poggendorff effect obtained for gapped shallow gradient stimuli is less than the effect obtained with the standard Poggendorff figure. Note that the contrast of the edges was maximal for this conventional stimulus. The implication is that the gapped stimuli are functionally equivalent to a lower than maximal contrast edge. That is, there exists a complete stimulus defined by a low contrast edge that induces the same sized Poggendorff Illusion as do the gapped stimuli. The purpose of the following experiment was to determine the effective contrast of the shallow gradient gapped bar and line stimuli.

Method

Subjects. Twenty-five first year University of Western Ontario Psychology students took part in the study to fulfill a course requirement. Eight were male and 17 were female. Ages ranged from 18 to 24. All had normal or corrected to normal vision. None had participated in any previous studies.

Stimuli. Six constant contrast bar and line stimuli were used to assess the effect of contrast on the Poggendorff Illusion. The vertical edges of these stimuli were physically complete from top to bottom. Contrast was defined by the formula: $(L_{background} - L_{stimulus})/(L_{background} + L_{stimulus})$, where L denotes luminance. The six contrasts were 91.51%, 7.57%, 4.41%, 3.27%, 2.16%, 0.88%. A transversal-only figure was also used, essentially a 0% contrast figure. Background luminance was 97 candelas.

A gapped bar and line stimulus were presented to the subjects in order to compare the size of the effect obtained with low contrast edges against that induced by the zero contrast contour. Gap angle for these two stimuli was 60 deg counterclockwise from horizontal and its size was 1.35 deg of visual angle. The contrast gradient of these stimuli was relatively shallow with a slope of 4.26 cd m^{-2} deg^{-1}. A ramp gradient was used in the present study to allow comparison between the current study and the results of the detection experiments reported in the previous chapter.

All stimuli were 1.97 deg of visual angle wide, centred on the display. The transversals were oriented at 65 deg from the horizontal. The 15 displays were presented six times each, resulting in a total of 90 trails. Each subject was run in a different random order of presentation. The transversals-only display and the highest contrast bar stimulus were used for practice.

Procedure. The method of adjustment was used: Subjects made six adjustments for each display. In other respects, the procedure was the same as that used in Experiments 3.2–3.4.

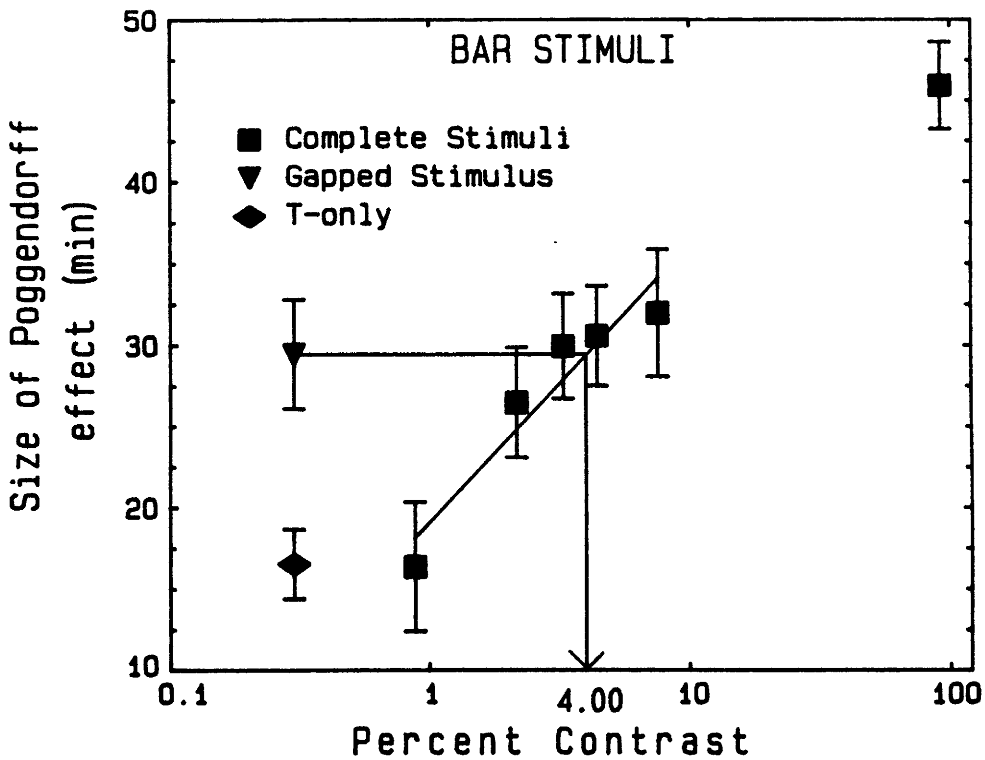

Figure 3.20: The Poggendorff effect as a function of contrast, bar stimuli (Experiment 3.6).

Results

The mean difference from true alignment, calculated across all subjects, is shown in Figures 3.20 and 3.21, for the bar and line stimuli, respectively. The x-axis in both figures denotes log percent contrast, and the y-axis, the size of the Poggendorff effect. The triangles show the effects obtained with the complete, constant contrast stimuli. The inverted triangle shows the effect obtained with the transversals-only stimulus,[3] and the square symbol, the effect induced by the gapped stimulus.

For both types the bar and lines stimuli, the results replicated those of Weintraub and Krantz (1971). The size of the Poggendorff Illusion de-

[3]Since log(0%) is undefined, this condition cannot be graphed on the scale of the x-axis. It was placed in the figure arbitrarily between 0.1 and 1.0 log unit.

creased as the contrast of the parallels decreased. The lowest contrast stimulus (0.88%) induced no greater effect than a transversals-only display. To provide an estimate of the effective contrast of the gapped stimuli, linear regressions were performed for each stimulus type separately. The size of the effect was regressed onto log percent contrast according to the procedure described by Sokal and Rohlf (1969, chapter 14) for multiple values of y (effect) at each level of x (contrast). Neither the 91.51% contrast nor the 0% contrast stimulus figures were included in the regression since the effect induced by the gapped stimuli, by inspection, lies somewhere in the 0.88–7.57% range, and the data are most linear in this range. A pilot study showed that the Poggendorff effect reached asymptote at approximately 20% contrast, suggesting that the effect is not linear beyond 20% contrast.

For the bar stimuli, the relationship between the expected size of the Poggendorff effect ($\hat{P}$) and $\log_{10}$ percent contrast (c) was found to be:

$$\hat{P} = 17.08c + 19.13$$

The 29.42′ effect induced by the gapped bar is estimated to correspond therefore to an effective contrast of 4.00%. Note that the effect was not significantly less than that induced by the complete 7.57% contrast bar, $t(24) = 1.33$.

For the line stimuli, the regression equation was

$$\hat{P} = 19.08c + 17.05$$

The estimated contrast for the gapped line stimulus was 3.03%. In this case, the effect induced by the gapped line stimulus was significantly less than that of the 7.57% contrast stimulus, $t(24) = 4.99, p < .001$, but failed to be different from the 4.41% contrast display, $t(24) = 1.15$. In sum, the data showed that gapped shallow gradient stimuli were functionally equivalent to complete stimuli of between three and four percent contrast.

General Discussion

The purpose of this set of experiments was to assess the functional role of smooth and abrupt edge endings. If smoothness increases the likelihood that a contour is interpolated through a gap, then that gap should function in a manner similar to the way a contour does, and, as the gap is made

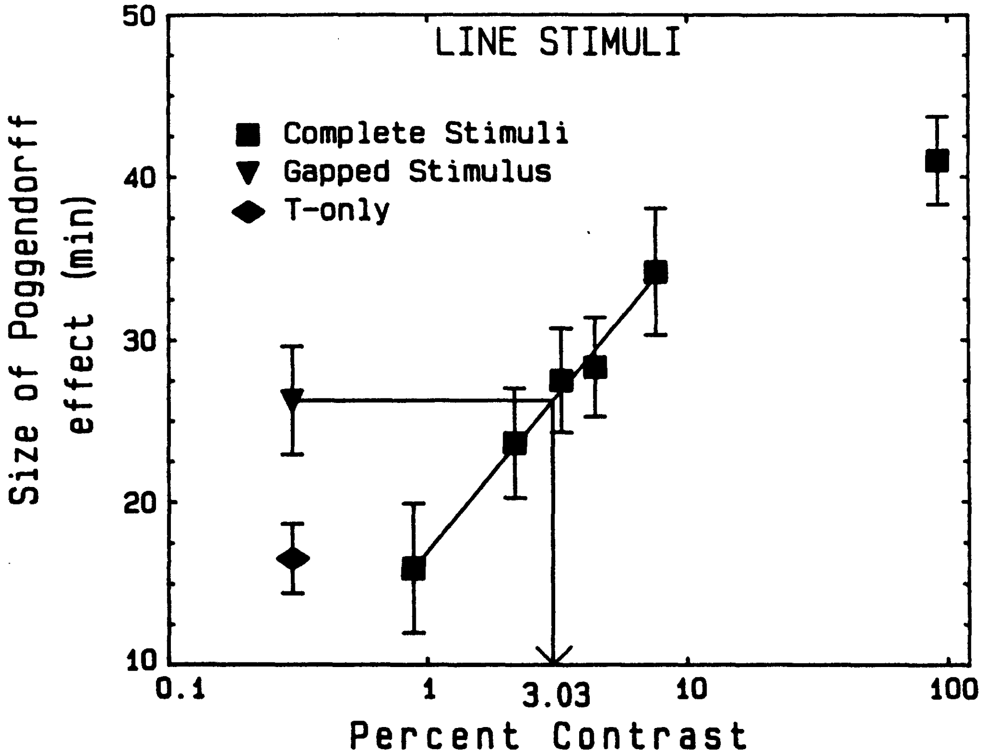

Figure 3.21: The Poggendorff effect as a function of contrast, line stimuli (Experiment 3.6).

more definite by making the endings more abrupt, then it should function less and less like a contour and more and more like a gap. Gapped stimuli consistently induced Poggendorff effects when they were defined by relatively shallow contrast gradients. These effects were greater than those obtained when subjects were presented with a transversals-only stimulus. The gapped stimulus did not behave, therefore, as if they were gapped, but as if there was a low contrast edge in a region where, in fact, there was no luminance difference. In addition, to a first approximation, as the gradient slope was increased, the effect decreased to that of a transversals-only stimulus.

With respect to the effects obtained when the gradient was steep and the edge endings were abrupt, two results are of note. First, the orientation of the gap relative to the angle of the transversals affected subjects responses even when there were definite contours framing the gap. For the framed bars, the

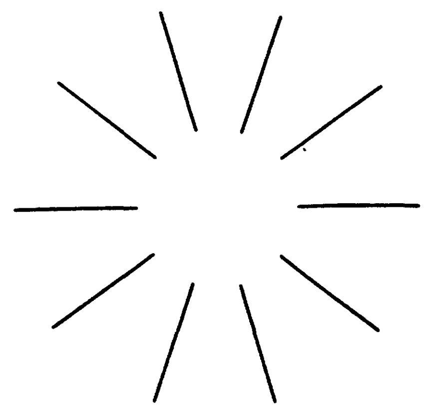

Figure 3.22: The Ehrenstein Illusion.

size of the Poggendorff was contingent upon the clarity of the gap—the clearer the gap, the less the misalignment. Second, for all three types of displays (framed bars, non-framed bars, and the line stimuli), as the orientation of the gap was made less different from the angle of the transversals, the size of the "Poggendorff" effect decreased to a point where it was equivalent to the effect induced by a transversals-only display, and in some instances even less. Thus, when the gradient was steep, the major factor determining subjects' responses was the orientation of the gap.

A possible explanation of this result is as follows. When the gradient is shallow, it is difficult to see exactly where the top and bottom of the zero contrast region are without careful scrutiny of the stimulus. It is these top and bottom edges that define the difference between gap and transversal orientation. When they are not easily seen, there is little or no such difference. In short, the gap is not a very good stimulus for a gap. On the other hand, when the gradient is steep, the top and bottom edges are distinct; the extent of the gap and its orientation are clearly visible. In this case, the gap is an entity, perceptually distinct from other parts of the stimulus, and possessing a definite orientation. Why it affected subjects' perception of alignment precisely the way it did is unknown. But, the fact that gap orientation did influence subjects' judgments argues that it is functioning as a gap. The visual system is not extending nor interpolating contours through it in this case; rather, the system is using the boundaries of the gap in the determination of the alignment of the transversals. In sum, when edge segments end abruptly, they are encoded as such and their ending influences the perception of the stimulus. Exactly how perception is modified can vary from situation to situation. For example, in some cases, an obscuring disk is inferred from the endings (see Figure 3.22). In the case of the Poggendorff Illusion, perceived alignment was altered.

In one of the experiments, no effect of gradient direction was found. Reversing and non-reversing gradients induced roughly the same sized Poggendorff Illusion. Recall that there was some evidence of a difference in gap detectability (Chapter 2) for reversing vs. non-reversing gradients. For reversing gradients, the gap was detected with greater ease, when a single bar was presented, and was judged more readily as "gapped" when compared to a complete bar.

This inconsistency is most likely due to a difference in task demands. In one case, subjects were asked to look for a gap. In the other, they were required to adjust a tilted line to apparent alignment. There is little reason to believe that subjects' criteria for these two tasks are the same (or even of the same type). Consider, for example, the finding of Experiment 3.6 that the "contrast" of the gap was effectively three to four percent. The data of the detection studies showed that a *real* one percent contrast edge was almost never seen as "gapped", in the no-cue condition. And, in fact, the gapped stimuli were rated as less "gapped" in appearance than the physically "gapped" stimulus. Obviously, the zero contrast region of the gapped stimuli did not have the *appearance* of three percent contrast edges, yet they effected the same consequence as a real three percent contrast edge, at least with respect to the Poggendorff.

This difference between the two experimental paradigms recalls the notion that perception is modular. At one level (detection), gaps bounded by smooth endings have the appearance of extremely low contrast edges. But, at another level, the level at which the Poggendorff Illusion occurs, they influence perception as if they were of greater contrast, i.e. they have the functional appearance of higher contrast edges. The factors that contribute to the percept at two given levels may well overlap, but it is doubtful that they are coextensive. Thus there is little reason to expect that the results of the two methodologies would agree at every point, and, specifically, that gradient direction is important for both the detection of gaps *and* the aligning of the transversals.

Nevertheless, it was expected that there would be some parallelism between the two different paradigms. If smooth edge endings do increase the likelihood that a contour completion mechanism is engaged, then it is expected that gaps bounded by smooth endings would elicit fewer "gapped" judgments and induce a Poggendorff Illusion. Abrupt endings are expected to evoke fewer "gapped" judgments and not induce a Poggendorff Illusion. These hypotheses were confirmed. Smooth endings did lead to contour interpolation on both counts. In the next chapter, the contour completion effects observed in the experiments are discussed in the broader context of contour perception.

Chapter 4

Related Phenomena

The results of the two sets of experiments have confirmed that the perception of a gap between two edges is determined, in part, by the manner in which those edges end. If the contrast gradually fades to zero, then there is an increased likelihood that the gap will not be seen as empty space, but as a continuation of the edges. Contours are interpolated under these conditions. Furthermore, the contours so interpolated are functionally similar to those defined by luminance differences, at least insofar as they induce a Poggendorff Illusion.

When the gap is signalled by luminance discontinuities, then it is readily seen as a gap, even to the extent of its size and orientation. It is not merely an area of constant luminance but is perceived as a channel separating the upper and lower halves of the stimulus displays. In the case of a Poggendorff stimulus, its orientation relative to the angle of the transversals modifies the perception of the alignment of those transversals. If the disparity between gap orientation and transversal angle is small, then perceived alignment is no worse than that found in a transversals-only display. As the disparity is increased, the size of the effect also increases. This suggests that the zero contrast region is a perceptual entity that affects the perception of the stimulus as a whole.

What kind of perceptual process is responsible for this completion effect? I will argue that it is an instance of contour interpolation. Recall from Chapter 1 that scene contours are not always projected in their entirety into the image. The corresponding edges sometimes end, and gaps occur among them. This was illustrated in Figure 1.1 where part of the occluding contour of the egg was missing. In this case, and others similar to it, it is desirable to fill in the gap since it does not correspond to an actual interruption of the underlying scene contour. The completion effects obtained in the experiments represent the output of a system designed to decide when gaps in edges ought to be filled in. The problem is solved, in part, by noting the manner in which the edges end. If they end smoothly, and there is no other information in the image to indicate that they truly end, then they are candidates for extension. If they end suddenly, or if there are cues that imply the ending is veridical, then that ending ought to be taken as given.

There exists the possibility that these completion effects are not instances of contour perception, but are the result of some other perceptual mechanism. In the following chapter, the contour completion effects found in the experiments are compared and contrasted with other completion phenomena. The chapter begins with a section on what contour completion definitely is not, in which it is compared to grating induction (cf. McCourt, 1982), phantom contours (cf. Tynan & Sekuler, 1975) and subjective contours (cf. Kanizsa, 1976). Next, a class of phenomena are considered that bear a striking resemblance to contour completion; effects such as the Craik-O'Brien Illusion (Ratliff, 1972). In this case, the similarities are noted, as well as certain differences that make it doubtful that contour completion is a "Craik-O'Brien Illusion". In the final section, we consider what contour completion is good for, which again argues in favour of its being an instance of contour perception.

4.1 What contour completion is not

4.1.1 Grating induction

The grating induction effect (McCourt, 1982) occurs when a homogeneous grey test field is placed perpendicular to, and overlaying a sine-wave luminance grating (see Figure 4.1). The appearance of the test field is that of a sine wave grating equal in spatial frequency to that of the inducing field, but of opposite phase. The areas of the test field adjacent to the dark bars of the surround appear lighter, and those adjacent to the light bars, darker.

One could argue that grating induction is a case of contour completion in that it causes the perception of "bars" within the test region, albeit of opposite sign to the "bars" in the inducing field. The induced bars align

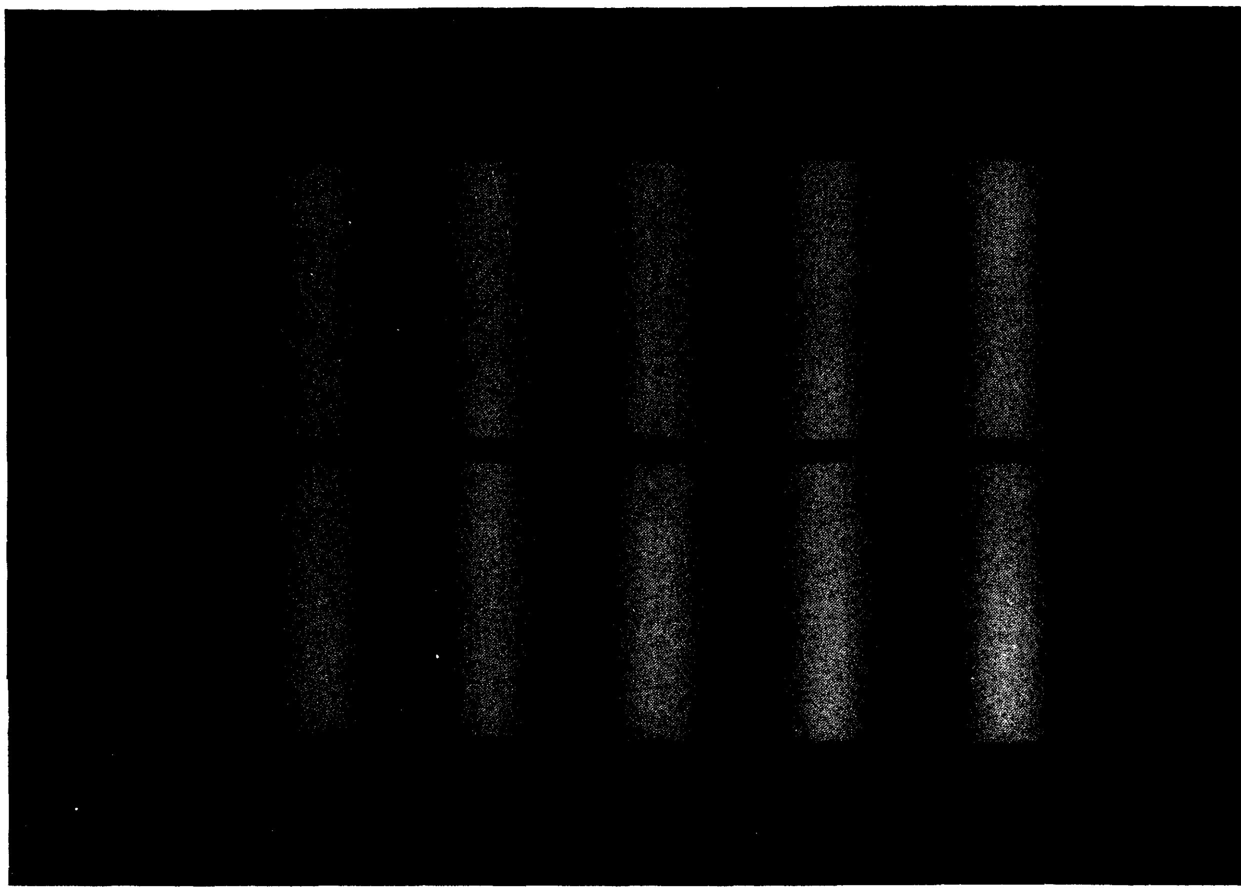

Figure 4.1: The grating induction effect (McCourt, 1982).

with the "bars" of the inducing grating such that, in a sense, contours are perceived running parallel to the induction pattern and through the test field. In other words, although the inducing grating is interrupted by the test field, the edges of the inducing pattern are perceived to continue through it. It is in that sense that the edges of the inducing grating are interpolated across a region where there are no physical luminance changes. Note further that the physical edges of the sinewave grating end suddenly where they intersect the test field. Thus, if grating induction were an instance of contour completion, then it would constitute a counterexample of the hypothesis that contour completion is inhibited by sudden endings.

Grating induction is probably not directly involved in contour perception, but is a function of mechanisms used for the determination of lightness, where "lightness" refers to the perceived luminance of a region of the stimulus. McCourt certainly investigated the grating induction effect under the assumption that he was studying a lightness phenomenon. Furthermore, in some ways, the effects he obtained resemble those of simultaneous contrast (cf. Heinemann, 1955, 1972) in which two test patches of identical luminance are seen as different when one is surrounded by a darker inducing field and the other by a lighter field. Two findings of McCourt's are discussed below in order to explicitly draw out the parallel between simultaneous contrast and grating induction.

Using a cancellation method, McCourt (1982) found that the effect is a lowpass phenomenon. Grating induction is greatest when the inducing field is of relatively low spatial frequency. The effect is attenuated at higher inducing frequencies, commencing at about three cycles/deg. A similar effect has been reported by Shapely (1984) using displays more closely resembling standard simultaneous contrast stimuli. He found an effect by placing two identical grey patches on either end of a shallow ramp luminance gradient. The grey patch on the dark end of the gradient appeared lighter than the patch on the light end. McCourt's and Shapley's research shows that low spatial frequencies affect perceived luminance: Gradual changes in the luminance of a surrounding field influence the lightness of the surrounded field.[1] McCourt, (1980) speculated that standard simultaneous contrast was a degenerate instance of grating induction. In this case, the inducing "grating"

[1] The effect can be seen in Figure 1.4 where the background (not the bars) is of constant luminance, yet the top appears brighter, and the bottom, darker.

is of the lowest possible spatial frequency: zero cycles/deg.

Second, McCourt (1982) found that for test fields greater than 0.6 deg in height, the effect was virtually nonexistent for induction fields whose frequency was greater than three cycles/deg. Foley and McCourt (1984) extended this result and found a complex relationship among induction frequency, test field height, and induction field height. Grating induction is a negative exponential function of these three factors. Induction is enhanced by an increase in induction field height, and diminished by an increase in the product of induction frequency and test field height. Thus, in order to obtain an effect with a large test field, the frequency of the inducing grating must be made fairly low and the size of the inducing field large. Note that standard simultaneous contrast is also dependent in a similar way upon the relative sizes of the test and induction fields. Simultaneous contrast is enhanced using smaller test fields and/or large induction fields (Heinemann, 1972).

Grating induction and contour completion

If grating induction is an instance of lightness perception, then it may be that the edges seen in the test field are merely a secondary effect based on the prior lightness computation. That is, given that the visual system has decided that two adjacent regions (of the test field) differ in luminance, it may infer further that there is an edge between the two regions. A difference in brightness, be it real or only apparent, implies that there is an edge marking the boundary between the lighter area and the darker.

Secondary perception of edges based on apparent luminance differences depends entirely on the luminance characteristics of the stimulus required to induce the primary brightness effect. This is not the type of information used in contour completion, where interpolation depends on properties of the endings of the edges, not specifically on luminance properties of the stimulus. This difference is reflected in the results of the experiments reported in Chapters 2 and 3 where the effects are seen to depend on factors that would not lead to grating induction. These are discussed below.

First, the contrast of the stimuli that induced contour completion were defined by relatively shallow contrast gradients. This meant that these stimuli were of fairly low overall contrast. For example, the *maximum* contrast of the bar defined by the (shallow) 4.47 cd m^{-2} deg^{-1} gradient was 12.25%, in the detection studies. In short, interpolation occurred for stimuli with lower average contrast and, as it were, was inhibited by higher contrast (steep gradient) stimuli. In comparison, simultaneous contrast is caused by high contrast inducing fields (Heinemann, 1972). Grating induction was produced also using relatively higher contrast "bars": McCourt (1982) used a contrast of 54%; Foley and McCourt (1984) used a 70% contrast grating. In their study of the effects of an interference grating on grating induction, McCourt and Foley (1985) used a 27% contrast inducing grating and a 27% contrast interfering grating.

Now, although the effects of the contrast of the inducing grating has not been specifically studied, grating induction is in all likelihood lessened by lower contrast inducing gratings. McCourt and Foley (1985) did vary contrast in one study to determine the threshold of inhibition by an interfering grating. They had found earlier that grating induction was diminished if a slightly higher frequency interference grating was added to the stimulus, and they wanted to determine the minimum contrast of the interfering grating needed to inhibit the grating induction effect. They found that, as the contrast of the interfering grating was increased, grating induction was decreased, suggesting that the greater the interference, the less visible the inducing grating, and the less the induction. That is, grating induction is probably not obtained when the inducing grating is of low contrast. By comparison, interpolation was obtained with the lower contrast stimuli in the experiments reported in Chapters 2 and 3, which suggests that contour completion is therefore not grating induction.

Second, grating induction was studied using "reversing gradients", meaning that the top of the inducing grating above the test field was in phase with the bottom. In the contour completion studies, some of the stimulus bars were defined by non-reversing gradients, in which the top of the bar was darker than the background, and the bottom, lighter. Recall that these stimuli generated as much contour completion as those involving reversing gradients. One suspects that grating induction would be severely impaired, if not eliminated, by counter-phase top and bottom inducing fields. The top would induce one pattern within the test field that would be countered by an opposite pattern induced by the lower field. Grating induction should be minimal in such a display; contour completion was unaffected.

Finally, recall that grating induction is a lowpass phenomenon. McCourt (1982) found that high frequency gratings did not induce the effect. Although the spatial frequency components of the stimuli used in the interpolation experiments are unknown, thin lines were found to induce a Poggendorff illusion in the experiments reported in Chapter 3. This suggests that contour completion between smooth endings does not depend on the spatial frequency of the inducing edges and, again that contour completion is not grating induction. The inducing stimuli can be large, low frequency bars or small, thin lines. However, before this can be concluded definitely, it would be necessary to determine the spatial frequency characteristics of the stimuli, and whether they play a role in the interpolation of the contours.

4.1.2 Phantom contours

Tynan and Sekuler (1975) described a completion effect, which they termed "phantom contours", that is in some ways similar to McCourt's grating induction. The stimulus in this case consisted of a vertical sinusoidal grating interrupted by a black horizontal strip. When the grating was made to drift horizontally across the display, a dim grating appeared within the black strip, which was in phase with the inducing grating, drifting in the same direction. When the surrounding grating was stationary, however, no phantom grating was seen.

Tynan and Sekuler reported the following properties of this completion effect. First, like grating induction, phantom contours are a lowpass phenomenon: As the spatial frequency of the induction grating was increased, the phantoms became more difficult to see. Second, when the black strip was rotated to the vertical, the phantom contours were no longer seen. Third, using a dichoptic presentation such that the top of the induction grating was presented to one eye and the bottom to the other eye, the phantoms remained equally vivid as compared to binocular viewing conditions. These latter two properties imply that phantom contours are the result of cortical processing; they are not a retinal phenomenon. Fourth, when square waves were used to induce the effect, the phantom contours took on the appearance of square waves.[2] Fifth, Tynan and Sekuler (1975) found that the phantom gratings were made more vivid when *lower* contrast inducing gratings were used.

One aspect of phantom contours that Tynan and Sekuler (1975) stressed was their strong dependence on movement. This has proved to be some-

[2] McCourt (1982) found that when he used a square-wave inducing grating, the induced pattern was neither a square-wave, nor a simple sinewave.

what debatable. Genter and Weisstein (1981) have shown that a very similar effect is induced by stationary gratings; however, the phantoms were seen only when the inducing grating was flickered, suggesting that the effect is dependent upon temporal modulation and not on motion per se. Unfortunately, this story is complicated by a finding of Gyoba (1983), who reported an effect obtained neither with motion nor flicker. He found "that stationary non-flickering gratings will also produce stationary phantoms ... if a high contrast grating is used" (Gyoba, 1983, p. 205).

What is the relationship, if any, between phantom contours and the contour interpolation effects induced by smooth edge endings? If phantom contours were entirely dependent upon some sort of temporal modulation, then there would appear to be no relationship. Contour completion did not depend on any temporal factors, but occurred in stationary stimuli. However, Gyoba's finding that phantoms can arise in stationary displays does suggest there may be some relationship between the two. Even so, it is likely that these are two distinct phenomena. In the first place, as was pointed out with respect to grating induction, contour completion was observed when the stimuli were of relatively low contrast. Gyoba (1983) reported using a 60% contrast grating in order induce stationary phantoms. Tynan and Sekuler (1975) did report that low contrast stimuli induced more vivid phantoms, but this occurred only when the inducing grating was moving. Thus, phantoms are seen with low contrast moving (flickering) gratings, or with high contrast stationary gratings. In either case, these effects are unlike those of contour interpolation which were obtained with low contrast stationary stimuli.

A second property of phantoms that reinforces the idea that they are unrelated to contour interpolation is the fourth property cited above: When a square wave inducing grating was used, the phantoms had a square wave appearance. Tynan and Sekuler (1975) also reported a "phantom dot" effect in which a blank rectangular region was surrounded by a drifting pattern of random dots. Observers reported seeing dots across the entire display, the dots in the blank region appearing dimmer. Weisstein, Maguire, and Berbaum (1977) investigated phantom contours using a "grating" made up of alternating black and patterned bars (see Figure 4.2). The patterned half of the grating consisted of white X's on a black background. When this grating was made to drift, the phantom contours took on the appearance of the inducing gratings. That is, a moving black-and-patterned grating appeared within the blank strip—the columns of X's were dimly but clearly

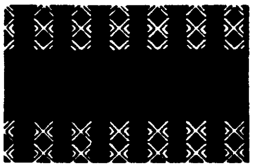

Figure 4.2: Phantom contour display used by Weisstein et al. (1977).

seen. These findings suggest that phantom contours are the result of a region based perceptual mechanism. Phantom contours result from the filling in of an entire region in terms of luminance and textural properties. Therefore, the contours do not appear to be a species of contour perception, but a secondary phenomenon arising from a region filling-in process. Where the visual system decides one region ends and another begins, that is where a contour is inferred.

4.1.3 Subjective contours

Contour completion and subjective contours are related in at least a functional sense: They are both concerned with the perception of contours given minimal information. However, it is doubtful that they are accomplished by the same mechanism. There is one point of intersection between the two that suggests that they are (almost literally) orthogonal to one another. With respect to contour completion, it was found that abruptly ending edges did not induce completion. Such endings constituted "news" to the effect that the edge should not be extrapolated. On the other hand, subjective contours are easily induced by sudden line endings, which also argues in favour of the notion that the lines themselves are not extended. The endings are used to infer another contour.

A sampling of subjective contour displays is shown in Figure 4.3, in which the subjective contour is induced by abrupt terminations. In Figure 4.3a, a vernier offset is seen as a pair of parallel horizontal lines. In 4.3b, the

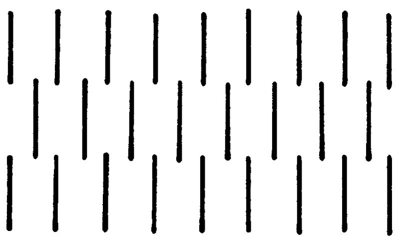

(a, from Ware, 1981)

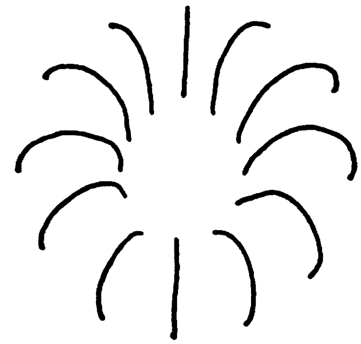

(b, from Brady & Grimson, 1981)

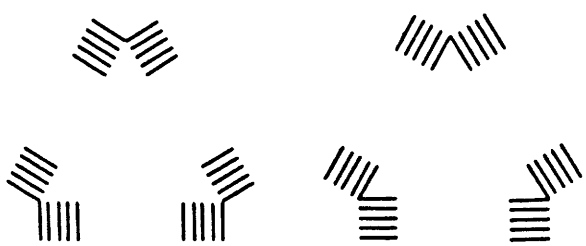

(c, from Frisby & Clatworthy, 1975)

Figure 4.3: Subjective contours induced by abrupt endings.

curved lines are seen as surface markings on an opaque but otherwise invisible doughnut. The terminations of the lines define locations where the surface markings pass behind the occluding boundaries of the doughnut. In this case the endings provide a cue to the occlusion, perhaps the only cue, and these occluding boundaries are seen even though no edge, corresponding to these boundaries, is present in the stimulus. Figures 4.3c and 4.3d show a pair of stimuli in which the inducing elements are placed in roughly the same positions. The difference lies in the orientation of the inducing elements; they are rotated 90 deg in one figure relative to the other, and this relative orientation affects the shape of the illusory countour. A triangle is perceived in 4.3c, and a Y-shape in 4.3d.

What these demonstrations show is that abrupt endings are (a) seen as endings; that is, the edge segments are not extended beyond their termination, and (b) the endings are used by the visual system to infer contours "orthogonal"[3] the ending. Contour interpolation is evident in these types of stimuli, however, the illusory contour is not a completion of the edge segments. The interpolated contour is one that intersects the edges and this intersection is cued by the abrupt terminations. Thus, part of the reason that subjective contours are perceived in these kinds of displays is because the edge endings are taken to be indicative of occlusion.

It is tempting to speculate that the subjective contours seen at abrupt endings could be annulled, or at least reduced, if the endings were made more smooth. No careful study was made of this, but there is some anecdotal evidence that when you juxtapose a number of vertical gapped bars involving smooth endings, a bistable percept results. If you see the gap, then you see a fuzzy, fog-like horizontal bar occluding the vertical bars. If you see the vertical bars as complete, then the subjective horizontal bar is less vivid. It is not clear whether both can be seen simultaneously. In order to make a more definite conclusion, this issue would have to be explored more systematically.

Subjective contours and lightness perception

There is an account of subjective contours, which enjoys some recognition in the literature, that constitutes an alternative explanation of why abrupt edge endings induce subjective contours. The impetus of this alternative is the fact that subjective contour displays give rise to a concomitant illusory brightness effect. For example, the subjective triangle in Figure 4.3c appears lighter than the surrounding background. Brigner and Gallagher (1974), Frisby and Clatworthy (1975), Jory and Day (1979), and Kennedy and Lee (1976) all have proposed that it is this induced brightness that leads to the perception of subjective contours. Specifically, they proposed that the inducing elements of subjective contour stimuli cause a perception of enhanced brightness near them (if the inducers are dark relative to the background), which the visual system then integrates. A large region of the stimulus is thereby seen as brighter and an intensity change (i.e. an edge) is seen at the appropriate locations. With respect to abrupt edge endings in particular, both Frisby and Clatworthy (1975) and Kennedy (Kennedy and Lee, 1976; Kennedy, 1978b)

[3]Not literally at 90 deg, but intersecting the edge (consider Figure 4.3b).

have proposed that a "button" of brightness is induced by the ends of thin dark lines, which, when there are a number of them (eg. Figure 4.3c), are joined together and define a larger region of apparent brightness.

On the face of it, this account of subjective contours differs from the one that focuses on the nature of the edge endings. Subjective contours, according to this account, are not seen because the visual system infers occlusion or interruption from a set of endings, but that these terminations induce an apparent intensity gradient.

There are reasons to doubt such an account of subjective contours. With the exception of showing that the ends of lines do induce simultaneous contrast (eg., Frisby & Clatworthy, 1975), most of the evidence in support of the brightness theory has consisted of a series of demonstrations in which the inducing elements are seen to cause both an apparent contrast and illusory contours; and, when no subjective contour is perceived, there is no corresponding lightness effect. This amounts to a strong correlation between the two phenomena, but that is not enough to conclude that brightness is primary in these situations. It could just as easily be the case that the perception of the contour is primary and that it, in turn, causes an illusory brightening of the perceived figure (cf. Bradley & Dumais, 1975; Kanizsa, 1976). What one would like to know is whether the correlation between contour and contrast is only a coincidence; whether the stimuli which induce one effect happen also to induce the other.

The way to discount the brightness explanation, then, is to break the correlation. A number of studies/demonstrations have shown that (a) illusory contours can be produced without concomitant induced brightness and (b) brightness effects arise in situations where no contours are induced. For example, Ware (1981) has shown, using magnitude estimation, that for a stimulus judged equally bright among three regions (Figure 4.3a), the subjective contour was judged only 16% less vivid than a comparable stimulus for which there was a distinct impression of contrast. That is, the brightness judgment decreased by a factor of 100%, whereas the contour vividness judgment changed only slightly.

Parks (1982) developed a version of a Kanizsa square in which the induced brightness effect was counteracted by a physical intensity gradient. Normally, such a subjective square is perceived as lighter than its background. Parks superimposed a luminance gradient on the stimulus such that the greater the radial distance from the centre, the brighter the stimulus. On average, the

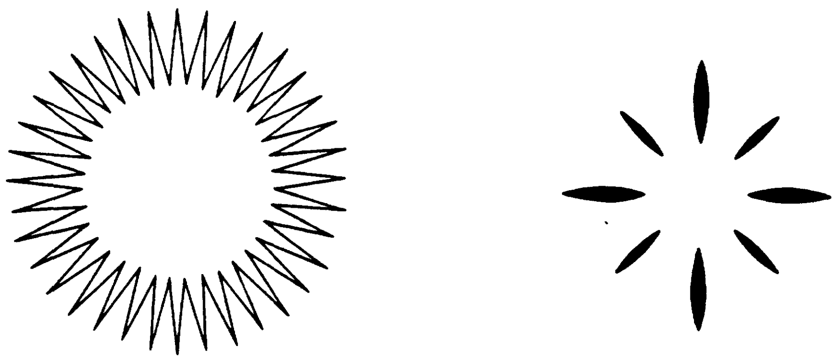

(a, from Kennedy, 1978a) (b, from Kennedy, 1981)

Figure 4.4: Star figure (a) and petals (b) that induce a brightness effect but no subjective contour.

display was physically darker inside the subjective contour than outside. Half of Parks' subjects reported that the subjective square appeared darker than the background. Also, they judged the same subjective square as lighter when it appeared on a homogeneous white background. The subjective contour was judged equally vivid in both conditions.

Kennedy (1976, 1978a, 1981) has published demonstrations that give rise to an apparent brightness effect but do not induce very strong subjective contours (see Figure 4.4). Petry, Harbeck, Conway and Levey (1983) have obtained judgments of illusory contour clarity and subjective contrast for these displays and found only a weak relationship between the two types of judgments. They concluded that apparent brightness and subjective contours were semi-independent phenomena.

With regard to abrupt endings, Prazdny (1983) constructed a display in which he attempted to nullify any brightness effect while maintaining the line ending cues. A facsimile is shown in Figure 4.5. The radial lines alternate between dark and light and lie upon a background whose luminance is the average of the lines. The endings of the dark lines should induce a local "button" of brightness, and the light lines a "button" of darkness. Together, these brightness effects should cancel and inhibit any global contrast phenomenon. Nevertheless, a definite circular subjective contour is seen in the figure, perpendicular to the line endings.

These experiments suggest that some stimulus conditions induce subjective contours without inducing a lightness effect; and, that there are figures

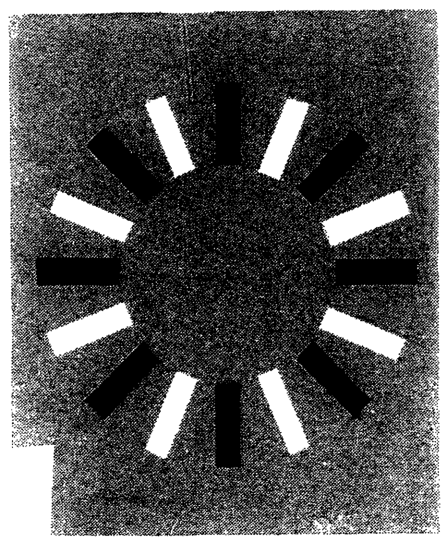

Figure 4.5: Radial lines of opposite contrast induce a subjective contour (from Prazdny, 1983).

that cause a brightness effect without leading to the perception of illusory contours. Given these findings, it is difficult to maintain that simultaneous contrast is the basis for the perception of subjective contours. In particular, faced with Prazdny's demonstration, the simultaneous contrast theory cannot explain why abrupt endings induce the perception of contours. It seems that regardless of any brightness effect induced by them, sudden terminations of edge segments are encoded by the visual system as such and used to infer other contours.

4.2 Gradient and edge effects

In this section, we consider a number of lightness illusions that depend on various kinds of luminance gradients. In some ways, these bear a strong resemblance to contour completion effects. There are significant differences, however, and it is unlikely that the two phenomena are the same.

An element that modifies the perceived lightness of a region is the presence or absence of an edge within the stimulus. The Benussi ring is a demonstration of this. A version of the Benussi ring is shown in Figure 4.6. This stimulus has the look of a standard simultaneous contrast display. Half of the grey ring is surrounded by a black rectangular region while the other half is bordered by a white region. Inspection of Figure 4.6a reveals some contrast effect: the half of the ring on the black background appears somewhat lighter than that on the white background. But the effect is not very substantial; certainly less than that usually found in standard simultaneous contrast stimuli. When, however, a thin black edge is added, as in Figure 4.6b, dividing the two halves of the ring, the apparent difference in lightness is enhanced. Thus, the Benussi ring shows that the presence of an edge enhances a perceived luminance difference, and, furthermore, the lack of the edge all but nullifies the effect.

Cornsweet (1970) reported a situation, diagrammed in Figure 4.7 which corresponds roughly to the Benussi ring shown in Figure 4.6a. The intensity distribution of this stimulus is such that the central region of the disk is physically brighter than the outside flanks. However, the change in luminance has been made very gradual such that no edge is seen. The luminance of the disk, when spun is shown in the upper curve of Figure 4.7, marked "L". The perception is of a grey disk constant in luminance across its face, as shown

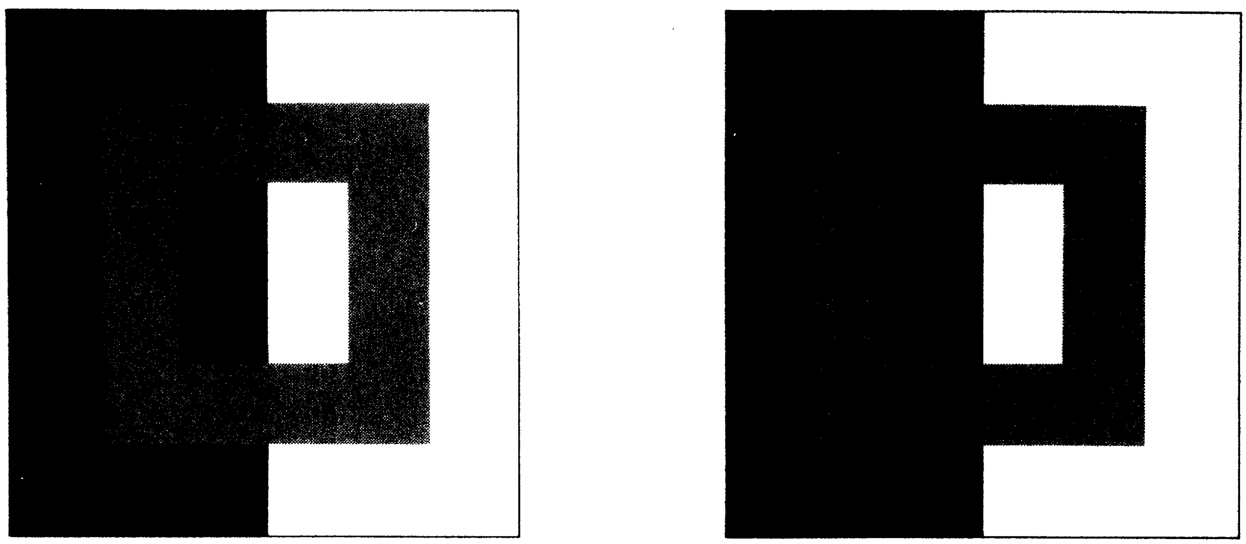

(a) (b)

Figure 4.6: The Benussi ring.

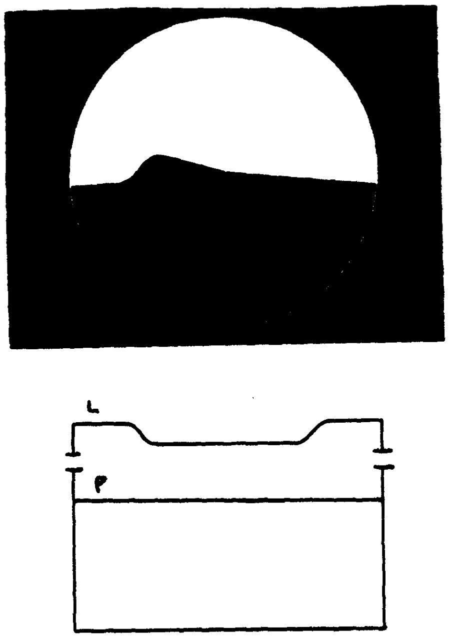

Figure 4.7: Schematic of the Cornsweet Illusion.

in the lower curve marked "P". Thus, even though there is a luminance difference between two regions, if the transition is smooth, then it is not seen.

Analogous to the Benussi ring, if a stimulus similar to Cornsweet's is constructed such that the luminance change between the two regions is sudden, then a luminance difference between them will be perceived. This occurs even under circumstances when there is no physical difference in luminance between the inner and outer parts of the disk. An example of this is the Craik-O'Brien effect, discussed by Ratliff (1972) and Cornsweet (1970). The stimulus properties of this display are diagrammed in Figure 4.8. The disk is constructed such that it is half black and half white over most of its face. A contour is placed at the midpoint between the outer circumference and centre of the disk. Except for this edge, the luminance of the disk, when spun, is constant across its face. This is diagrammed in the upper curve in Figure 4.8. The lower curve shows the perceived luminance: The lightness of the inner region bounded by the contour is seen as greater than that outside. The central region is lighter because the luminance is physically greater im-

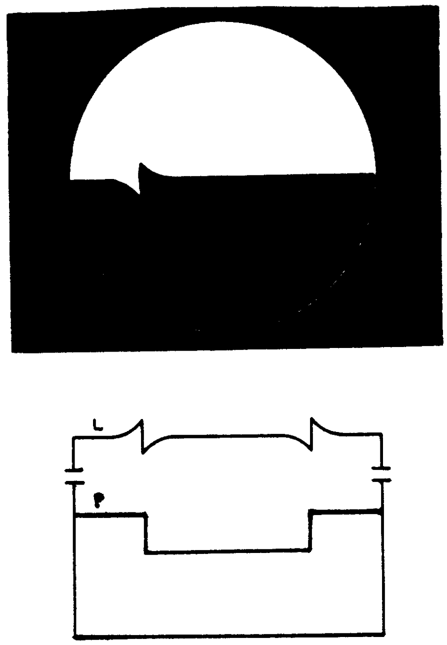

Figure 4.8: Schematic of the Craik-O'Brien effect.

mediately inside the contour and less on the outside. If this is reversed, then the central region appears darker (Ratliff, 1972). In a sense, the luminance immediately on each side of the edge "spreads" in either direction away from the edge until it encounters another edge.

These gradient and edge effects were noted by Land and McCann (1971; see also Horn, 1974) who used them as the basis for the retinex theory of lightness perception. They showed that if two regions differ only slightly in terms of their actual luminance, the perceived difference is enhanced by the placement of an edge between them. When the edge is removed, the apparent contrast is eliminated or greatly reduced. Land and McCann concluded that it is the edges that the visual system uses to compute the reflectance properties of the underlying surface patches, given their luminance profiles in the image. If there is a slow change in luminance between two regions, these are ignored unless they are cued by more abrupt changes.[4]

[4]Obviously, this cannot be the whole story. Recall that McCourt (1982) (grating induction) and Shapley (1984) showed that slow changes do affect lightness perception.

Similarities to contour completion

What the Benussi ring, Cornsweet-Craik-O'Brien effect, and Land and McCann's demonstration all show is that the lightness of a region is driven, sometimes, by the presence of luminance discontinuities. Furthermore, when two regions are separated by only a shallow luminance gradient, then the gradient is not seen as a change in luminance. The two regions, of physically different luminance, are perceived as possessing the same lightness. At an abstract level, this is similar to the idea that contour completion occurs when the edges end smoothly, but not when they end suddenly. The underlying principle is the same: Continuity (in luminance or edges) is implied by smooth variation, whereas discontinuities imply change.

The similarity may be stronger than that. Recall that in the gap detection experiments, the addition of cues encouraged the perception of a gap relative to a no-cue stimulus. The stimuli used in these studies (see Figure 2.1) look similar to the Benussi ring, and the completion effects follow a similar pattern to that of spreading lightness effects.

Like the Cornsweet Illusion, gap detection in the bar stimuli depended on the steepness of the contrast gradient. The failure to see the edge in the Cornsweet figure depends on how shallow the luminance gradient is. If it becomes too steep, then one sees the two regions distinctly. In a similar fashion, contour completion was observed when the gradient was relatively shallow. The percept is of a single continuous bar. In contrast, when the gradient was steep, the gap was easily detected. In this case, one sees three distinct regions, namely, the top of the bar, the gap, and the bottom of the bar.

The addition of the cues had an effect similar to the Craik-O'Brien Illusion. When an edge is placed on the disk, the spread of luminance is checked and is limited by that edge. In the detection experiments, when sharp luminance discontinuities (cues) were introduced, they delineated where one region ended and another began, and may have restricted the "spread" of the luminance gradient to the top and bottom halves of the stimuli. The gap was induced thereby to appear as a distinct region. Recall that this occurred even when the bar was complete: The cue masked the low contrast edge in the central part of the stimulus. In both cases, contour completion may have been hindered by the containment of the lightness spread to the top and bottom halves of the bar. These similarities between contour completion and the Cornsweet and Craik-O'Brien Illusions suggest that contour completion is an instance of lightness spread effect.

Differences from contour completion

First, let us be clear as to what such an effect would entail. These types of lightness illusions affect the perception of the luminance of a *region* of the stimulus. In the case of the Craik-O'Brien effect, for example, one does not only see that the disk is lighter immediately inside the edge, but the entire central portion of the disk is seen as brighter than the surrounding ring (see curve "P" in Figure 4.8). The regions on either side of the edge "assimilate" towards the sign of contrast on that side of the edge. The region on the brighter side looks brighter, and that on the darker, appears darker. Furthermore, the shallow gradient on either side of the edge is not seen. The edge effectively makes the luminance gradient invisible (Todorivić, 1987). With respect to the shallow gradient stimuli used in the detection experiments, which region was altered with respect to its apparent luminance? The only region that qualifies, if contour completion is an instance of the lightness spread, must be the central region of the bars, the part that was seen as gapped when the cues were added.

The question is, therefore, was the apparent *luminance* of the central portion of the bars affected by the cues, and/or the lack thereof? Certainly in the case of the stimuli defined by non-reversing gradients, the answer must be no. These were the stimuli that were darker than the background above the gap, and brighter below. Any spreading of the darker top portion of the bar would be cancelled by the spread of brightness induced by the lower half of the bar.[5] The gap would have to look either both brighter and darker than the background, which is impossible, or the same as the background. Phenomenologically, it is difficult to say what the luminance of the gap appears to be in the no-cue situation, when one perceives it as complete (see Figure 2.1). Whatever its apparent luminance, contours were interpolated in this instance, suggesting that the issue of lightness is orthogonal to whether contours complete across the gap.

As mentioned in Chapter 2, a lightness effect may have been induced in the case of the reversing gradient stimuli, in which the top and bottom of

[5] A similar sort of point was made by Prazdny (1983), mentioned above (see Figure 4.5).

the bar were both, on average, darker than the background. As remarked there, the gap in these stimuli appeared somewhat luminous and resembled a degraded Ehrenstein figure. But note, if there is a lightness effect in this case, it is the wrong kind. The Ehrenstein effect is a *contrast* effect, not an *assimilation* effect. The centre of the Ehrenstein figure appears lighter than its surround. The centre of the Craik-O'Brien disk has the same sign of contrast as the edge to which it is adjacent. If contour completion were a lightness spread effect, then the darkness of the top and bottom of the bars should have spread into the gap and it should therefore have looked darker than the background. If anything, however, the gap appeared brighter than background (and the top and bottom of the bars). This situation led to a contrast effect, not a lightness spread effect.

Is contour completion an instance of lightness spreading? In order to finally settle the question, experiments would have to be conducted in which one measured the apparent luminance of the central region of the stimuli. That is, subjects would not be asked whether they saw gaps or contours, but would be asked to judge the apparent brightness of the critical region, both with and without cues. The results would have to correlate perfectly with both the detection and Poggendorff experiments. In addition, if spreading effects are hindered only by the presence of some discontinuity, why did the constant background luminance not spread into the centre of the bar? That is, why did the gap not appear to be the same luminance as the background in the no cue condition? Perhaps it did, but, note, this was the stimulus that induced the fewest "gapped" judgments, and a substantial Poggendorff effect; so, if spreading of the surround into the gap occurred in this case, it did not appear to do so at the expense of contour completion. The point here is that lightness spread can occur in these stimuli in two different ways, only one of which could be conducive to contour completion. It is therefore premature to conclude that contour completion is a lightness spread effect; one would first have to determine the lightness of the centre of the stimuli.

4.3 What is contour completion good for?

Thus far the chapter has been concerned with what contour completion is not. This final section considers why contour completion is useful, and therefore, what functional role it plays in perception.

Edge detection has been regarded universally as an important initial activity of the visual system. Neurophysiologists (eg., Hubel & Wiesel, 1977; Schiller, Finlay, & Volman, 1976), psychophysicists (eg., Cornsweet, 1970; Ratliff, 1972), and researchers in machine vision (eg., Horn, 1973; Marr & Hildreth, 1980; Canny, 1983) have devoted considerable energy to this topic. Edge extraction is considered a useful goal for any visual system to pursue, and one reason frequently offered is that it is a way of recovering scene contours. The premise is that contours are projected into the image as changes in luminance (i.e. edges), and finding these is a first step in terms of representing the boundaries in a scene.

Locating luminance changes will not result in a complete representation of the scene contours, however. Edge detection failures occur when, for whatever reason, the scene contour is not fully projected into the image and there is no luminance change corresponding to it. Note that these failures are not rare occurrences. Several investigators have stated that edge detection techniques *typically* output disconnected edge segments that do not correctly reflect the connectedness of the contours in the corresponding scene (Ballard & Brown, 1982, chapter 4; Marr, 1975, 1982; Richards, Nishihara & Dawson, 1982; Ullman, 1976). The problem is made particularly acute by the fact that some gaps are veridical: Sometimes the gaps in an image are projections of gaps in scene contours; that is, sometimes boundaries do end. It is impossible to determine, based on intensity changes alone, whether an edge ending in an image is the projection of a true ending, or simply a poor projection. The problem is further compounded by the use of threshold techniques. Without a threshold, numerous false "noise" edges are detected. These are distinguished, however, by their low amplitude. By setting a threshold, they can be eliminated, but only at the expense of real, low amplitude edges. The number of false gaps (i.e. gaps that do not exist in the scene) is thereby increased. Edge detection is seen to have some shortcomings, then, as a method for finding scene contours. It sometimes fails to find an edge corresponding to a scene contour, and it cannot tell the difference between veridical and non-veridical endings.

That edge detection fails to recover scene contours should not be terribly surprising. Edge detection certainly cannot be necessary for contour perception since there are other sources of information in the stimulus to indicate the locations of the contours. Also, the human visual system is quite adept at capitalizing on these other sources and does not rely only upon the detection

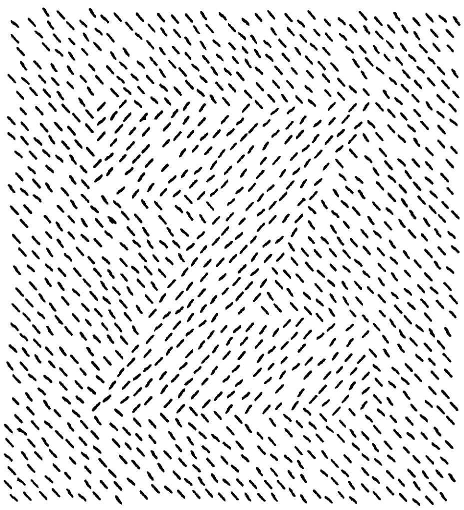

Figure 4.9: A Texture contour (from Riley, 1981).

of luminance changes.

One example of zero contrast contours are the disparity contours seen in random dot stereograms (Julesz, 1971). In this case, the perception of the contour depends on a perceived difference in depth. The images of the two eyes are compared, a correspondence is computed between them, and local disparity values are found. The stereoscopic contours mark locations of discontinuity in local disparity. Note that although changes in luminance in each of the individual images is involved in the computation of disparity contours, the contours themselves do *not* correspond to any luminance change.

Other situations that lead to the perception of contours, which are not defined by intensity changes, include contours induced by motion, and texture boundaries. Given two identical textural patterns (hence luminance distributions), if one region is moved against the background of the other, a boundary corresponding to the shape of the moving region is seen quite distinctly (Poggio & Reichardt, 1983). In this case, it is the relative motion that defines the perceived boundary.

An example of a texture contour is shown in Figure 4.9, taken from Riley (1984). Here, the intensity distribution within the Z-shaped region is identical to that without, yet a clear contour is seen bounding the "Z". The contour is defined by the change in orientation of the line segments comprising the figure. Riley reported that if this figure is filtered using a $\nabla^2 G$ operator (cf. Marr & Hildreth, 1980), at no scale is a zero-crossing found corresponding to the texture contour seen by human observers. That is, there is no luminance change corresponding to the texture boundary. He concluded that some other mechanism computed the boundary, using the relative differences in line segment orientation.

Here, then, are three examples of contour perception driven by something other than edge detection. In each case, the contour is inferred from some image property or properties—disparity, motion, and texture—not directly related to a luminance change in the image. By itself, edge detection would fail to find these contours.

Edge detection would also fail to recover the contour in situations such as that illustrated by the egg in Chapter 1 (Figure 1.1). Similarly, the gapped stimuli used in the experiments would all be "seen" as gapped by an edge detection algorithm. The gap in these displays was equal in luminance to the background—it was of *zero* contrast. An edge extraction scheme would (and should) fail to find an edge here since there was no edge in the image to be found. Nevertheless, the subjects perceived a contour in this region under certain conditions, specifically, when the gap was bounded by smooth endings and when there was no other cues to the ending. I believe that this is yet another case of the visual system using additional sources of information in the image to infer a contour where there is no change in luminance.

Edge detection could be improved, in terms of finding scene contours, if it noted the manner in which edges end. That is, in addition to localizing the obvious luminance changes in an image, some effort could be made to note whether the edge endings were abrupt or smooth. It would also be useful to look for corroborative information near the ending, like those explored in the detection experiments, to determine when an edge ending correctly depicts the underlying scene boundary. How the detection of these properties could be implemented and combined with the outputs of an edge detection algorithm is beyond the scope of this work. Nevertheless, if edge detection is used as a precursor of contour perception, then it is clear that the manner in which edges end must also be taken into account.

Chapter 5

Future Directions

In the previous chapter, the discussion focused on phenomena related to the contour completion effects found in Chapters 2 and 3. In the course of examining these relationships, additional experiments were proposed to investigate them further, and to tease apart the effects of various perceptual mechanisms that may underlie contour completion. There are, no doubt, other factors that affect contour completion and the present chapter reflects on a number of these.

Distance: One feature of the gap, which is probably important, is its size. Interpolation should occur more often when the edge endings are closer to one another. This is because the further apart two edge segments are in the image, the greater the likelihood that they are the projection of two different scene boundaries, in which case it would be an error to interpolate between them. In addition, it is doubtful that a single boundary would possess both of the following properties: (1) its projection extends over a large distance in the image, and (2) it is not registered by any change in the image over that length. If a contour is of considerable length, then, if anything, it will vanish for only small sections, and appear in patches in the image. Thus, it is undesirable to join edges that are separated by large distances for fear of inferring incorrectly that they are part of the same contour.

Distance is a factor in the case of subjective contours. Dumais and Bradley (1976) found that subjective contours were rated more vivid as the distance among the contour inducing elements was decreased. Also, the study of the effects of subjective contours on the Poggendorff Illusion (Experiment 3.1) showed that larger effects were obtained when the distance between certain of the inducing elements was small. A similar effect is predicted for gaps bounded by smooth endings. As the size of the zero contrast region is made smaller, the vividness of the interpolated contours should increase.

The effect of distance on interpolation could be investigated using the Poggendorff illusion and/or the gap detection paradigm. In the case of gap detection, the proportion of "gapped" judgments would be expected to decrease as the size of the zero contrast region was made smaller. With regard to the Poggendorff, the size of the illusion should increase as the distance between the endings decreased.

Shape: Another property that bears consideration is the shape of the interpolated contours. All of the figures used in the experiments were constructed from straight edges and it is relatively easy to tell how the edge segments should be connected. Curvilinear edge segments present a greater problem since there are numerous ways in which curved edges can be connected. The simplest method, for example, is to connect their end points with a straight line. However, in situations where the zero contrast region is bounded by curved edges, such as the egg in Chapter 1 (Figures 1.1), the visual system does not employ this method. Instead, it extends the edges in a manner that preserves their curved nature. Note that a similar phenomenon is observed in the case of subjective contours—one can, for example, construct a curved Kanizsa triangle (cf. Kanizsa 1976). On what basis does the visual system choose the path along which it will interpolate a (curved) contour? A smoothness constraint may supply the answer: Unless there is some indication in the image that the direction of the contour changes significantly, then the shape of the interpolated contour should not deviate much from that of the ending edges. A minimal curvature principle may apply to this situation (Ullman, 1976) in which the curvature at the endings constrain the shape of the interpolated contour in such a way as to minimize its curvature. A "kink" in the interpolated contour is to be avoided, and is allowed only if a smooth shape cannot be interpolated.

Locus of processing: Another question concerns where, in the visual system, the completion of the contour is achieved. It is doubtful that it has a retinal locus since it is difficult to imagine what retinal mechanism has the capability to effect the interpolation. The receptive field would have to be relatively large in order to span the 1.35 deg gap used in the experiments.

One way to determine whether the retina is the locus of the effect is to bypass it using a dichoptic presentation. The top half of the stimulus could be presented to one eye, and the bottom half to the other. In this case, no gap is present in either retinal image alone, but exists only at the point where information from the two eyes is combined. Replication of the effects observed in Chapters 2 and 3 under dichoptic presentation conditions would provide evidence relevant to the hypothesis that interpolation is done somewhere beyond the retina.

Surface interpolation: A related issue is which module, functionally defined, is responsible for the completion. The discussion has assumed that that a contour perception device is responsible. There is, however, a possibility that contour completion is based on a three dimensional interpretation of the image. In this regard, Brady and Grimson (1981) argued that subjective contours are a by-product of the visual system's attempt to construct a viewer centred three dimensional representation of the image. They proposed that when presented with *any* stimulus, the visual system constructs a three dimensional interpretation that is consistent with the features in the image, if such an interpretation is possible. Subjective surfaces are interpolated among the edges, regions, and "blobs" in the image, using these as constraints on the interpolation process. Subjective contours are simply the locations where there is a discontinuity in interpolated surface depth. For example, the subjective doughnut in Figure 4.3 is perceived because (a) the lines are seen as 3-d curves, and (b) their endings are consistent with an occlusion interpretation. The subjective contours are the occluding boundaries of the doughnut, marking where it ends and its background begins.[1]

Brady and Grimson's theory may account for the interpolation effects observed at smoothly ending edges. Most of the stimuli used in the demonstrations and the experiments can be given a three dimensional interpretation. Figure 1.1, for example, is a shaded egg shaped object. The stimulus bars defined by non-reversed luminance gradients have the appearance of curved shaded strips. Those defined by reversed gradients look like strips with glare across their centres. The only stimuli for which it is questionable that there is a three dimensional interpretation are the line stimuli used in the Poggendorff experiments. Even so, one can imagine that these lines correspond to

[1]The main problem with this theory is that, at present, it does not specify the nature of the features which initiate the interpolation mechanism. What distinguishes a stimulus interpreted as two dimensional from one inducing a three dimensional interpretation?

the shading and highlights along a pair of thin objects.

In order to research this question further, it would be necessary to introduce surface information in such a way that it conflicts with the contour information. Is it possible to design a stimulus in which they are inconsistent? Consider the following situation: Construct the stimulus such that the top of the bar has zero disparity along its entire length, whereas the bottom has some constant disparity value greater than zero such that the bottom half of the bar is seen as, say, above the background. The top half would appear flush with the background. What would the apparent depth of the gap be in such a situation? That depends, again, on whether the gap is bounded by abrupt vs. smooth endings. When the edges end suddenly, there is no inconsistency between the contour information and the surface information as they both change abruptly at the gap. The contours end, and the disparity shifts to some non-zero value. The perception of the stimulus would be of two bars, the bottom one seen in depth relative to the top. The gap would appear to lie in the same depth plane as the background, below the bottom half of the bar.

When the edges end smoothly, however, there is a conflict between disparity and contour information. The disparity still changes suddenly at the gap, but the edges gradually vanish. If surface interpretation supersedes contour interpolation, then the prediction is that the discontinuity in depth would hinder contour interpolation and the percept would be qualitatively similar to that obtained with abrupt endings. However, if a contour completion mechanism is engaged regardless of the surface information, then the vertical contours of the top and bottom sections of the bar would be joined together. In fact, I suspect, the visual system would probably compromise between the two sources of information. The edge segments would be joined, but the gap would appear to change depth from top to bottom. At the top, it would be of the same depth as the top half of the bar, and as it approached the bottom, gradually change depth to that of the bottom half of the bar.

Smooth endings: Finally, there is a question of how reliable smooth endings are. The reader may have been misled into thinking that smooth edge endings arise only when a scene contour is poorly projected into the image; that smooth endings never result from the projection of some smoothly ending scene boundary. This would be convenient from the point of view of contour interpolation since, if it were the case, smooth endings would constitute a signature that an edge's ending does not accurately reflect the

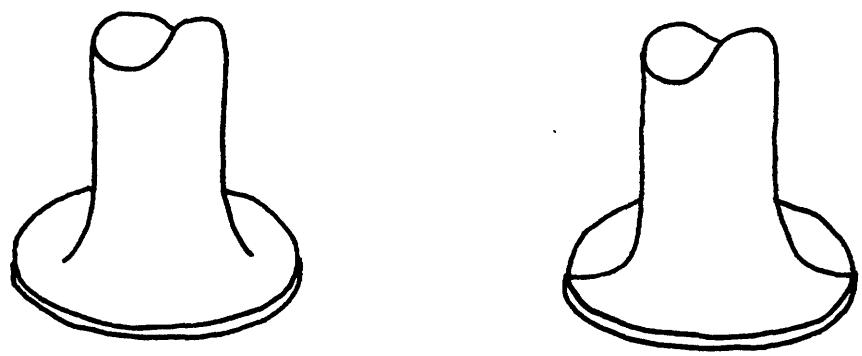

(a) (b)

Figure 5.1: Example of smoothly ending scene boundaries (from Koenderink & van Doorn, 1982).

properties of the scene. Unfortunately, the situation is not that simple. Some occluding boundaries do in fact end smoothly. An example, taken from Koenderink and van Doorn (1982), is diagrammed in Figure 5.1a. The vertical occluding boundaries of the pillar smoothly join its flared base and this is reflected in the image as a gradual decrease in the contrast of the corresponding edges. An edge that ends smoothly can indicate, therefore, a truly ending contour.

It should be noted that this is a *physical* description of the scene and image. It is an open question, *psychophysically* speaking, whether people correctly perceive edges like Figure 5.1a as ending. In point of fact, Koenderink and van Doorn (1982) implied that artists erroneously extend such contours until they meet the circular occluding contour of the base, as in Figure 5.1b. Obviously, this is not conclusive evidence that the contours are actually seen as continuing, but it is suggestive. Note specifically that the smoothly ending edges are extended until they meet some other definite contour, that is, that interpolation (if it occurs here) is halted by another "cue" in the image.

Evidence that the incorrect drawing of the pillar in Figure 5.1b does reflect the perception of the vertical contours is given by one final demonstration,

Figure 5.2: Smooth endings and cues.

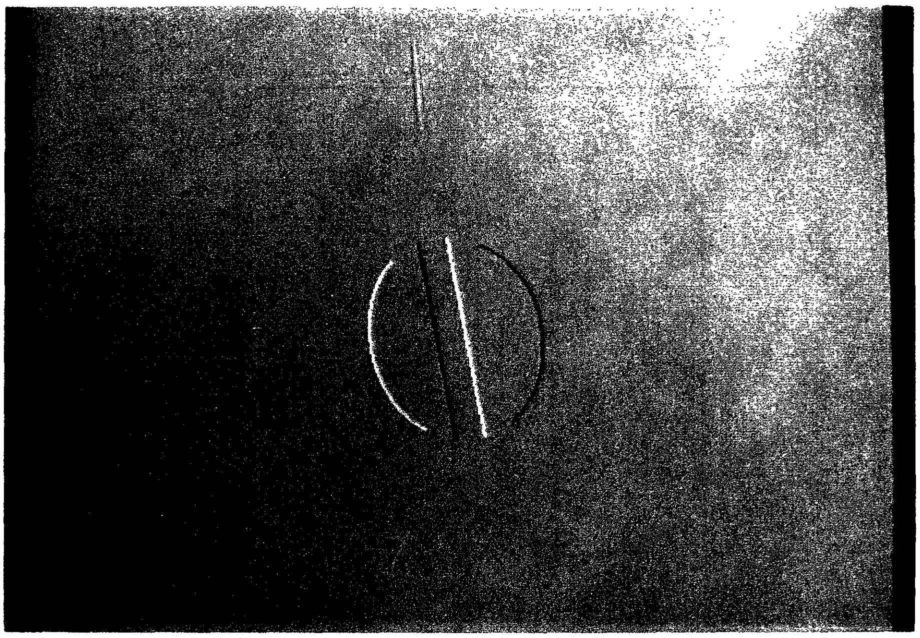

Figure 5.3: Locations of the endings.

shown in Figure 5.2 which, again, argues that the visual system extends edges that end smoothly. The figure consists of a pair of smoothly ending curved edges separated by a pair of straight vertical lines. The impression is of two embossed semicircular shapes separated by a vertical channel. The important point of this demonstration is that the curved edges do *not* physically intersect the vertical lines, but end some distance apart from them. The perception of the curved lines is, however, that they extend to, and end at the vertical contours. The gap between the curved and straight edges is shown in Figure 5.3, a high contrast version of Figure 5.2. This demonstration suggests that smooth endings will always be extended over short distances, unless and until there is information to the contrary.

Concluding Remarks. These, then, are other factors whose role in contour completion deserves further exploration. No doubt others will be discovered in the process of doing the research. In addition, there is the issue of how all these variables interact with other sources of contour information. Some discussion was devoted to this topic above with respect to the relative contributions of edge endings and stereo disparity. One must consider also how, for example, motion and texture fit into the picture. All of these factors must be investigated before contour perception is understood.

Chapter 6

Coda: Two Views of Terminators

This thesis began as a series of studies concerning terminator detection. Specifically, the issue was whether terminators were detected in parallel or required a serial search. Obviously, the research followed a different course while continuing to focus on "terminators". Below is a retrospection of how this came about. None of the arguments are intended as conclusive; rather they constitute a cursory reflection on my current world view.

The terminator detection research was done in the context of Julesz' (1981) "texton" theory and Triesman's (cf. Triesman & Gelade, 1980) concept of "basic features". According to this view, visual perception is a series of inferences where each step uses the output of the previous and provides input for the next. Tracing the process backwards far enough leads eventually to the most primitive inputs: the ones that initiate the process in the first place. The distinguishing characteristic of basic features is that they are seen effortlessly; they do not require an act of attention, but are detected in parallel.

For the most part, I am partial to this view. Indeed, the core concept, of vision as an inference process, played a significant role in the development of the research I eventually undertook. The dissatisfaction with the basic feature view arose in connection with the notion that a basic feature is detected in parallel: I began to doubt that being detected in parallel meant a great deal.

There was first a suspicion that being detected in parallel did not constitute a *necessary* condition for something to qualify as a basic feature. There is nothing in the concept "basic feature" that requires parallel search. All that "basic" entails is that the feature is as primitive as can be, and that it cannot be reduced to the output of some more basic process. By itself, this cannot mandate that basic features be detected in parallel.

Looked at another way, is it really inconceivable that there are some basic features that have to be picked up in serial? In other words, might there be basic features that, by their very nature, require serial detection? Surely the possibility exists. What is required to establish parallelism as a necessary condition is to show that no basic feature is in point of fact ever picked up in serial. That might turn out to be true, but it strikes me as unlikely. The visual system appears to be more flexible than that and is willing to do just about anything, if it proved necessary. If serial search were required in some cases, the visual system is not above giving up on parallelism.

Parallelism could be salvaged by retreating to a claim of sufficiency: Parallel search is *sufficient* for something being a basic feature. This is a much weaker claim. On this view, the world is no longer cleanly divided into basic and non-basic features by the litmus test of parallelism. Given only sufficiency, the set of basic features will include a potential subset of serially detected basic features. That entails that if a feature were found to be detected in serial, there is no way of knowing whether it is basic or not. Parallelism alone will no longer generate a complete list of the basic features.

In addition, there is the possibility that a non-basic feature can be detected in parallel. Schneider and Schiffrin (1977), for example, have argued that with enough practice, certain mental acts can become "automatic" and will cease to require attention to control the process. If this is true, parallel detection will fail even as a sufficient condition.

I think that what is really at the root of the parallel/serial distinction is speed. If there are basic features, then it behooves the organism to get them processed fast. The natural world is not a static display; it changes rapidly from moment to moment. If you are going to use the information in the retinal mosaic at any given instant, you had better get a handle on it before it is replaced by something else. To be sure, parallel detection appears to have the upper hand here, since it is a fast way of acquiring information. But, if speed is what is essential, it does not really matter if you detect the information in parallel, as long as a fast serial detector will do the job.

Faced with these doubts I began to ask myself if there was a more fruitful conception of "basic feature". If the view of perception as a sequence of levels is taken seriously, then it follows that something is relatively basic if it can be demonstrated that it is used by a succeeding level of processing. The key word here is "used"—what defines whether something is relatively basic is whether it is useful to some later stage in perception. For example, consider the following episode: Julesz and Spivack (1967) claimed that the terminators in random line stereograms were sufficient to produce stereopsis. In other words, the disparity between the two sets of terminators provided all the necessary information. Nishihara and Poggio (1982) countered this claim by showing that the low spatial frequencies in random line stereograms were sufficient for stereopsis (terminators are considered high frequency information). Furthermore, they went on to show that if the low frequency disparity information was eliminated, and all the disparity information was indeed carried by the high frequency terminators, then human subjects could no longer achieve stereoscopic perception. It appears that stereo depth perception depends on disparity information carried by the low spatial frequencies.

The moral of the above is that by doing experiments of this kind, one can determine what sorts of features are "basic" relative to a particular perceptual process. Notice that the question of whether terminators (or for that matter, low spatial frequencies) were found in parallel was beside the point. In their attempt to determine what was used for stereo matching, Julesz and Spivack (1967), and Nishihara and Poggio (1982) modified the input to eliminate certain features while retaining others. In the end, if the results of a parallel search study were at loggerheads with their results, I suspect that the parallel search data would be judged less important.

I have attempted to apply this functional approach to terminators. Instead of asking whether they are detected in parallel or serial, the studies were designed to determine if terminators were used for anything. A parallel search paradigm will show that they are detected in parallel, but it will not tell you what terminators are *for*. Answering this latter question requires a different kind of analysis. You must first specify what a terminator is; that is, you must first characterize what information is being used by the system. I have chosen to characterize terminators as the locations where an intensity edge ends and the contrast of the edge goes to zero. How might that be informative? Initially, the intuition is that an ending is the projection of the termination of a scene boundary. But that need not be the case since edges end in, broadly speaking, two different ways. They either end suddenly, or they gradually fade away. Consideration of how the two types of endings arise, given that the edges are the projections of scene boundaries, leads to insight regarding how endings might be informative.

The difference in the two views resides in their focus. One view is concerned with the functional role various image properties play in the construction of perceptions. And, if lucky, it may be possible to determine that the detection of some image property is *necessary* input, in which case that property can be declared a basic feature. The other view may tell us what the set of basic features are, but it will not go beyond that. Ultimately, the functional role of these features will have to be assessed anyway. For these reasons, the functional exploration of perception tells us, I believe, much more than a "basic feature" approach does.

References

Ballard, D. H., & Brown, C. M. (1982) *Computer Vision.* Englewood Cliffs: Prentice-Hall Inc.

Bradley, D. R., & Dumais, S. T. (1975) Ambiguous cognitive contours. *Nature, 257*, 582–584.

Brady, M., & Grimson, W. E. L. (1981) The perception of subjective surfaces. AI Memo 666, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Brigner, W. L., & Gallagher, M. B. (1974) Subjective contours: Apparent depth or simultaneous brightness contrast? *Perceptual and Motor Skills, 38*, 1047–1053.

Carterette, E. C. (1984) A computer program for estimating thresholds by means of Wetherill's up-down transformed response rule (UDTR). *The Journal of Auditory Research, 24*, 191–204.

Canny, J. F. (1983) Finding edges and lines in images. AI-TR 720, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Cornsweet, T. N. (1970) *Visual Perception.* New York: Academic Press.

Day, R. H. (1973) The oblique line illusion: The Poggendorff effect without parallels. *Quarterly Journal of Experimental Psychology, 25*, 535–541.

Day, R. H., Dickinson, R. G., & Jory, M. K. (1977) The Poggendorff illusion with subjective contours. *Quarterly Journal of Experimental Psychology, 29*, 219–226.

Dumais, S. T., & Bradley, D. R. (1976) The effects of illumination level and retinal size on the apparent strength of subjective contours. *Perception and Psychophysics, 19*, 339–345.

Foley, J. M., & McCourt, M. E. (1985) Visual grating induction. *Journal of the Optical Society of America (A), 2*, 1220–1230.

Frisby, J. P., & Clatworthy, J. L. (1975) Illusory contours: Curious cases of simultaneous contrast? *Perception, 4*, 349–357.

Genter, C. R., & Weisstein, N. (1981) Flickering phantoms: A motion illusion without motion. *Vision Research, 21*, 963–966.

Gibson, J. J. (1979) *The Ecological Approach to Visual Perception.* Boston: Houghton Mifflin.

Goldstein, M. B., & Weintraub, D. J. (1972) The parallel-less Poggendorff: Virtual contours put the illusion down but not out. *Perception and Psychophysics, 11*, 353–355.

Gregory, R. L. (1972) Cognitive contours. *Nature, 238*, 51–52.

Grimson, W. E. L. (1981a) The implicit constraints of the primal sketch. AI Memo 663, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Grimson, W. E. L. (1981b) *From Images to Surfaces: A Computational Study of the Human Early Visual System.* Cambridge, MA: MIT Press.

Guilford, J. P. (1954) *Psychometric Methods.* New York: McGraw-Hill.

Gyoba, J. (1983) Stationary phantoms: A completion effect without motion and flicker. *Vision Research, 23*, 205–211.

Heinemann, E. G. (1955) Simultaneous brightness induction as a function of inducing and test field luminance. *Journal of Experimental Psychology, 50*, 89–96.

Heinemann, E. G. (1972) Simultaneous brightness induction. In *Handbook of Sensory Physiology, VII/4, Visual Psychophysics*, D. Jameson (Ed). New York: Springer-Verlag.

Helmholtz, H. von (1962, original publication 1866) *Treatise on Physiological Optics, III.* J. P. C. Southall (Ed). New York: Dover.

Hildreth, E. C. (1980) Implementation of a theory of edge detection. AI-TR 579, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Horn, B. K. P. (1973) The Binford-Horn line finder. AI Memo 285, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Horn, B. K. P. (1974) Determining lightness from an image. *Computer Graphics and Image Processing, 3*, 277–299.

Hubel, D. H., & Wiesel, T. N. (1977) Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society of London (B), 198*, 1–59.

Jory, M. K., & Day, R. H. (1979) The relationship between brightness contrast and illusory contours. *Perception, 8*, 3–9.

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

Julesz, B. (1981) Textons, the elements of texture perception, and their interactions. *Nature, 290*, 91–97.

Julesz, B., & Spivack, J. G. (1967) Stereopsis based on verneir acuity cues alone. *Science, 157*, 563–565.

Kanisza, G. (1976) Subjective Contours. *Scientific American, 234*, 48–64.

Kennedy, J. M. (1976) Sun figure: An illusory diffuse contour resulting from an arrangement of dots. *Perception, 7*, 479–481.

Kennedy, J. M. (1978a) Illusory contours not due to completion. *Perception, 7*, 187–189.

Kennedy, J. M. (1978b) Illusory contours and the ends of lines. *Perception, 7*, 605–607.

Kennedy, J. M. (1981) Illusory brightness and the ends of petals: Change in brightness without aid of stratification or assimilation effects. *Perception, 10*, 583–585.

Kennedy, J. M., & Lee, H. (1976) A figure density hypothesis and illusory contour brightness. *Perception, 5*, 387–392.

Kirk, R. E. (1968) *Experimental Design: Procedures for the Behavioral Sciences.* Belmont, CA: Brooks/Cole.

Koenderink, J. J., & van Doorn, A. J. (1982) The shape of smooth objects and the way contours end. *Perception, 11*, 129–137.

Land, E. H., & McCann, J. J. (1971) Lightness and retinex theory. *Journal of the Optical Society of America, 61*, 1–11.

Marr, D. (1976) Early processing of visual information. *Philosophical Transactions of the Royal Society of London (B), 275*, 483–524.

Marr, D. (1982) *Vision.* San Francisco: W. H. Freeman & Company.

Marr, D., & Hildreth, E. (1980) Theory of edge detection. *Procceedings of the Royal Society of London (B), 275*, 483–524.

Meyer, G. E., & Garges, C. (1979) Subjective contours and the Poggendorff Illusion. *Perception and Psychophysics, 26*, 302–304.

McCourt, M. E. (1982) A spatial frequency dependent grating-induction effect. *Vision Research, 22*, 119–134.

McCourt, M. E., & Foley, J. M. (1985) Spatial frequency interference on grating-induction. *Vision Research, 25*, 1507–1518.

Mosteller, F. (1951) Remarks on the method of paired comparisons: III. A test of significance for paired comparisons when equal standard deviations and equal correlations are assumed. *Psychometrika, 16*, 207–218.

Nishihara, H. K., & Poggio, T. (1982) Hidden cues in random line stereograms. *Nature, 300*, 347–349.

Parks, T. E. (1982) Illusory contours: On the efficacy of their need for expression. *Perception and Psychophysics, 32*, 286–289.

Petry, S., Harbeck, A., Conway, J., & Levey, J. (1983) Stimulus determinants of brightness and distinctness of subjective contours. *Perception and Psychophysics, 34*, 169–174.

Prazdny, K. (1983) Illusory contours are not caused by simultaneous brightness contrast. *Perception and Psychophysics, 34*, 403–404.

Pylyshyn, Z. W. (1984) *Computation and Cognition: Towards a Foundation for Cognitive Science.* Cambridge MA: MIT Press, A Bradford Book.

Ratliff, F. (1972) Contour and contrast. *Scientific American, 226*, 90–101.

Reichardt, W., & Poggio, T. (1979) Figure-ground discrimination by relative movement in the visual system of the fly. *Biological Cybernetics, 35*, 81–100.

Richards, W. A., Nishihara, H. K., & Dawson, B. (1982) CARTOON: A biologically motivated edge detection algorithm. AI Memo 668, Artificial Intelligence Laboratory, Massachusetts Institutue of Technolgy.

Riley, M. D. (1981) The representation of image texture. AI-TR 649, Artificial Intelligence Laboratory, Massachusetts Institute of Technology.

Schiller, P. H., Finlay, B. L., & Volman, S. F. (1976) Quantitative studies of single cell properties in monkey striate cortex. I. Spatiotemporal organization of receptive fields. *Journal of Neurophysiology, 39*, 1288–1319.

Schneider, W., & Shiffrin, R. M. (1977) Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review, 84*, 1–66.

Shapley, R. (1984) The importance of contrast for the activity of single neurons, the VEP and perception. *Vision Research, 26*, 45–61.

Sokal, R. R., & Rohlf, F. J. (1969) *Biometry.* San Francisco: Freeman.

Spillmann, L., Fuld, K., & Gerrits, H. J. M. (1976) Brightness contrast in the Ehrenstein illusion. *Vision Research, 16*, 713–719.

Todorivić, D. (1987) The Craik-O'Brien-Cornsweet effect: New varieties and their theoretical implications. *Perception and Psychophysics, 42*, 545–560.

Triesman, A. M. & Gelade, G. (1980) A feature-integration theory of attention. *Cognitive Psychology, 12*, 97–136.

Tynan, P., & Sekular, R. (1975) Moving visual phantoms: A new contour completion effect. *Science, 188*, 951–952.

Ullman, S. (1976) Filling in the gaps: The shape of subjective contours and a model for their generation. *Biological Cybernetics, 25*, 1–6.

Ware, C. (1981) Subjective contours independent of subjective brightness. *Perception and Psychophysics, 29*, 500–504.

Weintraub, D. J. & Krantz, D. H. (1971) The Poggendorff illusion: Amputations, rotations, and other perturbations. *Perception and Psychophysics, 10*, 257–264.

Weisstein, N., Maguire, W., & Berbaum, K. (1977) A phantom-motion aftereffect. *Science, 198*, 955–957.