REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
31 August 1996

3. REPORT TYPE AND DATES COVERED
Final Technical 1 Jan 1992 - 30 June 1996

4. TITLE AND SUBTITLE
Visual Encoding of Spatial Relations

5. FUNDING NUMBERS
G F49620-92-J-0114
PR 2313
TA AS
61102F

6. AUTHOR(S)
Christina A. Burbeck

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of North Carolina at Chapel Hill
Office of Contracts and Grants
440 W. Franklin Street
Chapel Hill, NC 27599-1350

8. PERFORMING ORGANIZATION REPORT NUMBER
AFOSR-TR-96
0502

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AFOSR/NL
110 Duncan Avenue, Suite B115
Bolling AFB DC 20332-0001

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release;
distribution unlimited.

12b. DISTRIBUTION CODE

19961016 145

13. ABSTRACT *(Maximum 200 words)*

The two major findings from this work were the contribution of scaled boundary linking in shape and size perception and the existence of statistical representations for sets of objects. Linking boundaries of simple spatial regions at a spatial resolution proportional to the region's width yields the scaling of size judgment with size and provides a scale-invariant representation of simple object shape. This approach to image analysis has been adopted for the analysis of medical images. The idea that the visual system creates a statistical representation arose from findings that observers represent the mean value in a set with high precision but retain almost no information about the individual items in the set.

DTIC QUALITY INSPECTED 4

14. SUBJECT TERMS
Human vision, spatial vision, shape perception, localizaiton

15. NUMBER OF PAGES
26

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

29 AUG 1996

Visual Encoding of Spatial Relations

Christina A. Burbeck
Department of Computer Science
Sitterson Hall
University of North Carolina
Chapel Hill, NC 27599-3175

31 August 1996

Final Technical Report : 1 Jan 1992 - 30 June 1996

Prepared for
Air Force Office of Scientific Research/NL
110 Duncan Avenue, Suite B115
Bolling AFB, DC 20332-0001

Dr. John F. Tangney, Program Manager

Abstract

The two major findings from this work were the contribution of scaled boundary linking in shape and size perception and the existence of statistical representations for sets of objects. Linking boundaries of simple spatial regions at a spatial resolution proportional to the region's width yields the scaling of size judgment with size and provides a scale-invariant representation of simple object shape. This approach to image analysis has been adopted for the analysis of medical images. The idea that the visual system creates a statistical representation arose from findings that observers represent the mean value in a set with high precision but retain almost no information about the individual items in the set.

Objectives

[copied from the original grant proposal]

Our research focuses on understanding the visual processes that are responsible for encoding the sizes of objects. This work includes study of the basic properties of the size-encoding process and study of how this process interacts with other aspects of spatial vision. Of particular interest is the perceptual linking of regions of the image that appears to occur prior to the judgment of size. We have found that the accuracy of rapid size judgments depends on the similarity and location of background objects presented simultaneously. We infer from this that the ability of an observer to make a rapid and accurate size judgment depends on his initial perceptual organization of the image. Investigation of this organizational process has become an essential component of our research effort, being important in its own right and also helping to place the size-encoding process within the larger structure of visual processing as a whole.

We propose to investigate three aspects of the size-encoding process:

1) the properties of the process specifically devoted to encoding precise distances in the fronto-parallel plane.

2) **the interrelationship between the size-encoding process and object representation (i.e., representation of specific regions as belonging to a single object).**

3) **the relationship between the size-encoding process and the perceived spatial layout of a complex scene.**

I. Scientific Findings

We made substantial progress on all three goals, developing a model of shape representation that simultaneously accounted for the precision of size judgments and provided a means of segregating the scene into meaningful regions. Our interest in the relationship of size judgments to the perceived spatial layout of complex scenes led us to new ideas on how multiple similar objects are represented.

A. From Spatial Relations To Objects

1) Encoding Spatial Relations

For several years our research focused on the perception of the most simple spatial relations: the x-y distance between two locations. We showed the robustness of this percept across surface characteristics of the targets (Burbeck, 1987), exposure duration (Burbeck, 1986; Burbeck & Yap, 1990b), contrast (Burbeck, 1987), and retinal eccentricity (Burbeck & Yap, 1990c). We also explored the effects of non-target objects on the perception of the separation between a pair of targets. We found that at brief durations thresholds can be elevated (Burbeck, 1992; Burbeck & Yap, 1990a) or perceived size altered (Burbeck, 1993) by the presence of lines flanking the target lines (i.e., by distracters of a particular type). At longer durations (400-500 ms), these context effects were substantially diminished.

An important aspect of the finding that perceived size is altered by the presence of a flanking line was that the range of locations over which the flanking line has its effect depends on the target separation. Fig. 1 shows the stimulus and typical results. (Details and complete data are given in Appendix A: "Scaled Position Integration Areas: Accounting for Weber's Law for Separation".) The abscissa is the distance between the flanking line (the topmost line in the 3-line configuration) and the top target line (the middle line in the 3-line configuration). The ordinate is the change in the perceived size of the target separation that results from addition of the flanking line. The effect of this third line was modeled as the product of the distance to the flanking line and a Gaussian centered on the top target line. The standard deviation of the Gaussian and a single scale factor (for each observer and target separation) were free parameters. The standard deviation of the Gaussian increased systematically with the target separation. Further the rate of increase matched the rate of increase of the separation discrimination threshold over the range of separations tested. We inferred from this that the area over which position information is being integrated scales with the distance being spanned and that this scaling is sufficient to account for the increase in threshold with increasing size: Weber's law for size.

a)

b)

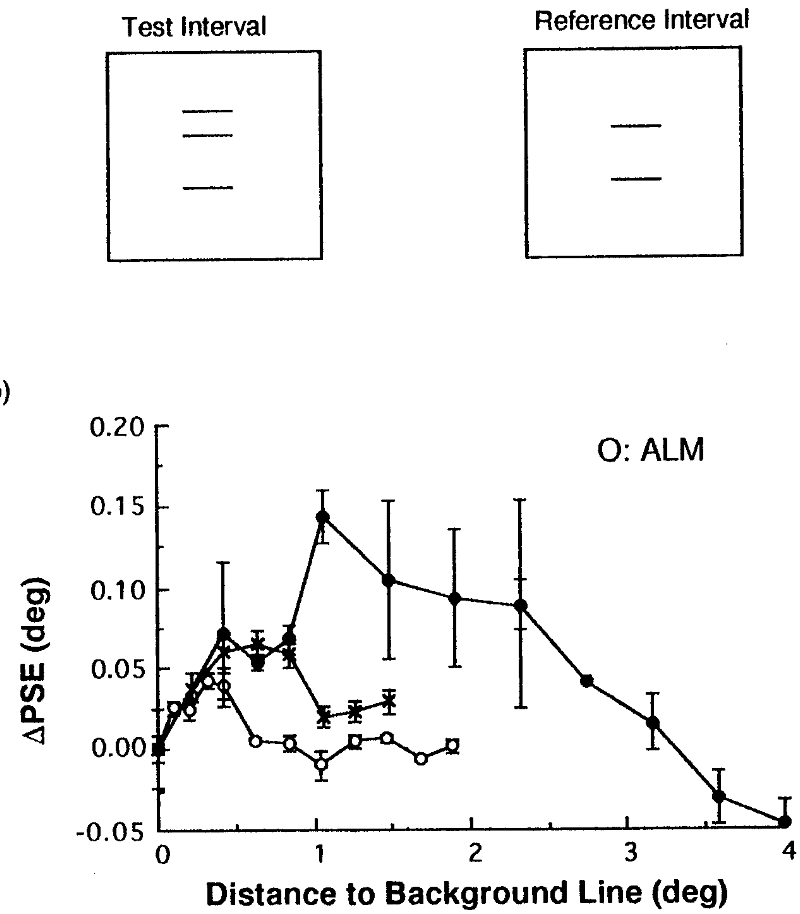

Fig. 1 a) Stimulus configuration used in Burbeck & Hadden (Burbeck, 1993). The test and reference intervals were 100 ms in duration, and each interval was terminated by the presentation of a masking stimulus. b) Typical results obtained from this experiment. ΔPSE is the increase in the perceived target separation for the stimuli in the test interval relative to the reference interval. The distance to the background line is the distance between the top two lines in the test interval. Filled circles, 3.0° mean target separation; crosses, 1.5° mean target separation; open circles, 0.75° mean target separation.

Additional studies in our laboratory showed that this effect was not due to luminance integration: a black flanking line had an almost identical effect. We also found that the integration area is circular: measurements made with a pair of target dots (see Fig. 2) and a pair of flanking dots, whose separation was varied, showed the lateral extent of the relevant area.

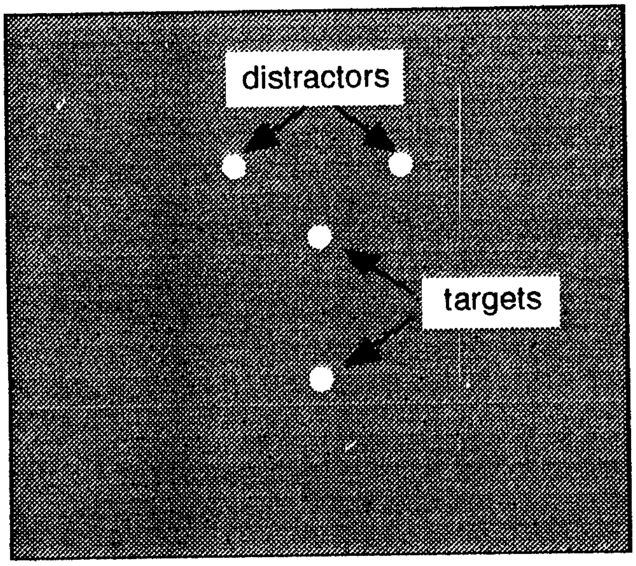

Fig. 2 Stimulus configuration used to determine the lateral spatial extent of the region contributing to the perceived location of the upper target.

Prior findings from our laboratory and others (Burbeck, 1987; Burbeck, 1988; Toet & Koenderink, 1988) showing that size judgments are independent of surface properties had already led to a two-stage model of size processing in which the second stage operates on the rectified response of the first stage. This idea was first proposed in 1987 (Burbeck, 1987) and has since been extended (Hess & Badcock, 1995). This idea has also been adapted to motion perception (Sperling, 1990).

Our flanking line results suggest that the second stage of processing consists of position integration areas that scale with the distance being spanned. Data obtained with intervening distracters (Burbeck, 1992; Burbeck & Yap, 1990a) suggest that this second stage does not consist of connected receptive fields, but rather linked, disconnected receptive fields. This disconnection renders the process insensitive to features between the targets and has the advantage that it enables the visual system to link different types of boundaries, e.g., a texture boundary with a luminance boundary. This proved to be an important concept in our work on shape representation.

2) The Core Theory of Shape Representation

Simultaneously with the latter part of the work described above, the PI began collaborating with Stephen Pizer, Kenan Professor of Computer Science at UNC, on the

2) The Core Theory of Shape Representation

Simultaneously with the latter part of the work described above, the PI began collaborating with Stephen Pizer, Kenan Professor of Computer Science at UNC, on the problem of shape representation. The central idea of their joint effort was that the 2D shape of a region is, in essence, the spatial relations between the boundaries of the region. In particular, we hypothesized that opposite boundaries (those most nearly parallel) should be explicitly linked in a representation of shape. The nature of the linking was strongly directed by two essential factors: the need for at least approximate zoom invariance and conformation with Weber's law for size. These requirements were both met by using boundary apertures whose size scales with the distance being judged, as suggested by our experimental work on size judgments (Burbeck, 1993) and theoretical work on zoom invariance (Pizer, Burbeck, Coggins, Fritsch, & Morse, 1994).

The mechanism for boundary linking that we proposed was as follows (see Fig. 3). Boundariness detectors vote for a middle and a width in a direction perpendicular to their preferred orientation at a distance from themselves that is proportional to their scale. We call the results of such voting by all boundariness detectors, "medialness". It is a value determined by the likelihood that a given location corresponds to the center of a region of a given width (with that width being proportional to the scale of the voting boundariness detector). When the location and scale of the votes from different boundariness detectors agree, there is an enhanced response. The locus of this enhanced response is the medial representation that we use. We call it the core of the region.

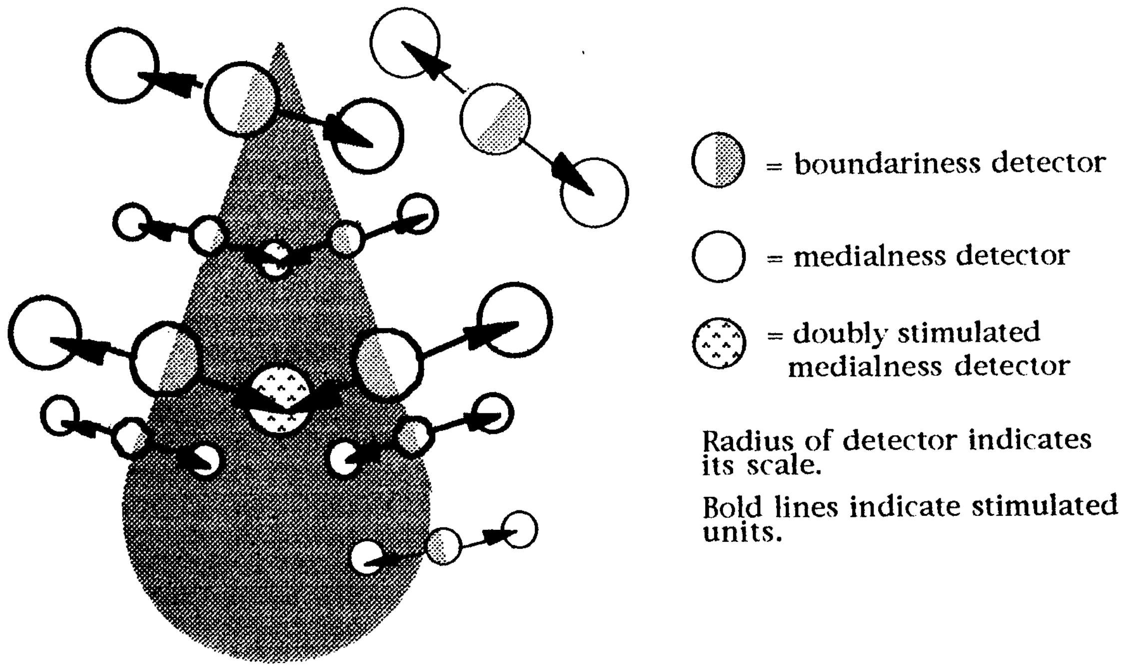

Fig. 3 (Fig. 4 from (Burbeck & Pizer, 1995)) Boundariness detectors combining (or failing to combine) to produce strong medialness on a teardrop-shaped region.

The core theory is described more fully in the Burbeck & Pizer's (1995) paper, "Object Representation by Cores: Identifying and Representing Primitive Spatial Regions". Briefly, a core is a trace in the 3-dimensional space consisting of two spatial dimensions (our model applies to 2-D profiles only) and a scale dimension. This trace represents the location and local width of a region, both of which are acquired and encoded with a spatial resolution proportional to local width — in accordance with zoom invariance and Weber's law for size. This medial axis model has the unusual advantage that it does not require that the boundaries be found first. Instead, the graded responses of boundariness detectors produce a graded medialness response, which has a ridge, which is the core. The intrinsic spatial scale of the core (which is determined by the spatial scales of the linked boundariness detectors) establishes a level of resolution for the representation that is appropriate for the object, rather than that resolution being determined either by fiat or by the resolution limits of the system. Further, the two-sidedness of the core gives it a robustness across weaknesses and gaps in the boundary that can be used, in turn, to aid in locating the boundary more precisely if desired (Pizer, 1994).

The core model avoids many other problems inherent in previous medial models (Blum, 1973; Blum & Nagel, 1978; Brady, 1983; Leyton, 1992; Marr & Nishihara, 1978).

The scaling — with object width — of the area over which boundary information is gathered makes the core insensitive to noise, gaps, protrusions, etc. at the boundary whenever their spatial scale is small relative to the local object width. (Other medial models create axis branches for every blip at the boundary.) The core model has been implemented computationally and is being applied in several medical contexts under NIH sponsorship (through a program grant, "Medical Image Presentation", Stephen Pizer, PI, from the National Cancer Institute).

3) Experimental Tests of the Core Theory: The Linking at Scale and Medial Representation Hypotheses

Although experimental results led us to the primary assumption of the theory, that boundaries are linked at a spatial scale determined by the local width of the object, alternative models can always be devised. Therefore it was important to test the specific predictions of our model after we had developed it. We also were interested in testing our hypothesis that boundary linking is done through a medial locus, and we had no specific data on that subject. Consequently, we conducted a series of experiments designed to measure the scale at which boundaries are linked, to determine whether that scale depends on object width, and to test the idea that the middle of an object carries shape information.

In our first experiments, the task was bisection. The stimuli were sinusoidally-edge-modulated objects of the type shown in Fig. 4. A small probe dot was placed near the center of an edge-modulated stimulus — in line with a peak of the modulation — and the observer was asked to report whether the dot was to the left or right of the local center of the object, where local meant the center along a line through the probe dot. Exposure duration was 500 ms, giving the observers time to foveate the probe dot. From the observers' judgments of the perceived center (testing at two vertical locations, one in line with a leftward peak and one in line with a rightward peak), we inferred the perceived modulation of the center of the object. Measurements were made with a range of edge modulation frequencies, from 1 to 32 cycles/object for 4° long objects, with two edge modulation amplitudes and two object widths. The idea was that if boundaries were linked at a spatial scale determined by the object width, then a) the perceived middle should become increasingly straight as edge frequency increased — reflecting the fact that the boundary information is acquired over a spatially extended area, and b) for a given edge frequency, a wider object should look straighter than a narrower object. Both results were found. Some sample data are shown in Fig. 4. The full results were given in "Linking Object Boundaries at Scale: A Common Mechanism for Size and Shape Judgments" by Burbeck, Pizer, Morse, Ariely, Zauberman and Rolland (1996).

a)

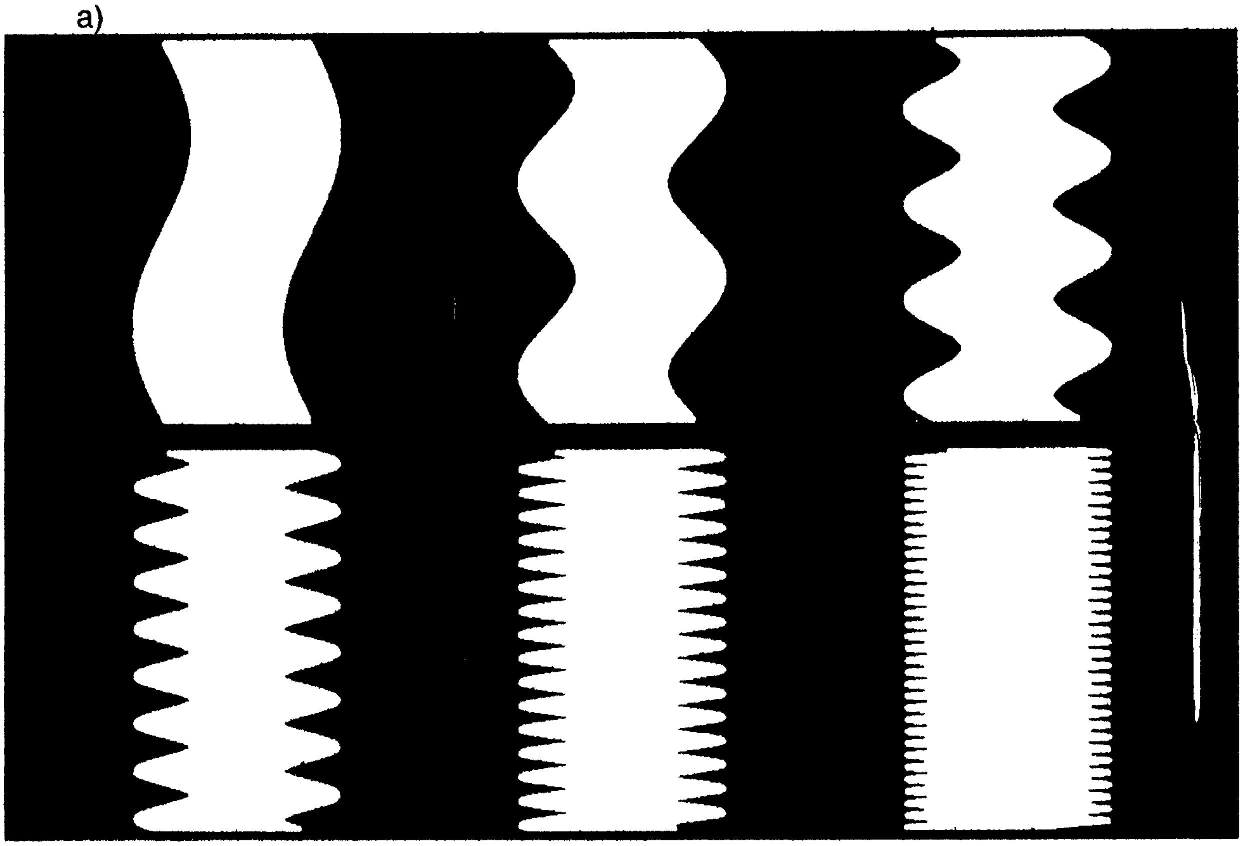

b)

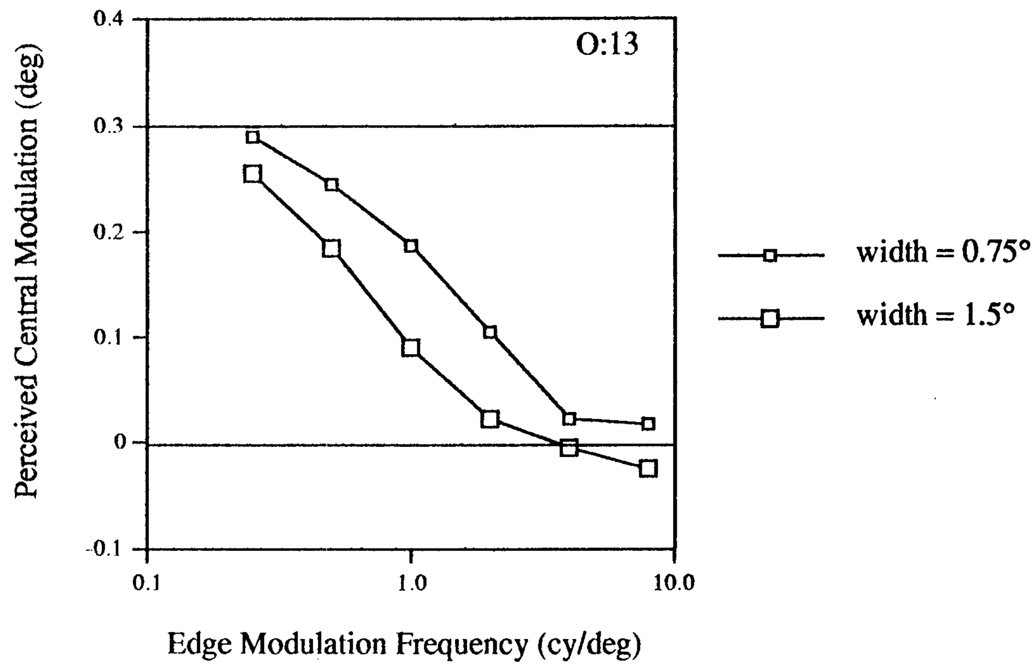

Fig. 4 a) Two of the sinusoidally-edge-modulated stimuli used in (Burbeck, Pizer, Morse, Ariely, Zauberman, & Rolland, 1996). The edge frequency here is 0.75 cy/deg. The edge modulation amplitude is 0.3°. b) Perceived central modulation of the wiggly-edged object as a function of the frequency of the edge modulation and the width of the stimulus. The effect of the edge modulation depends on the width of the object.

In addition to these primary findings, we found that, for these stimuli, the area over which boundary information is gathered does not scale exactly with object width: a doubling of the object width causes less than a doubling of the relevant integration area. Significantly, the increase in boundary integration area scales exactly with the bisection thresholds, supporting the idea that a common mechanism — scaling of the relevant boundary integration areas with object width — is responsible for both events. This result leant firm support to our basic theoretical premise: boundary linking at scale. It also supported the idea that shape information is carried in the medial locus: as the object became perceptually straighter, its medial locus did also.

We also tested the idea of a medial representation using a different paradigm, with the same sinusoidally-edge-modulated stimuli. The task was orientation discrimination. The idea was that edge-modulated objects with straight centers (created by modulating the two edges in opposite phase), would have a clearer perceived orientation than would edge-modulated objects with wiggly centers (i.e., with in-phase edge modulation). Our results were interesting and surprising.

The straight objects did indeed yield more consistent results: for the straight objects perceived orientation was constant across edge modulation frequency (for a given observer and object width) whereas for the wiggly-centered ones the error in the perceived orientation varied across these parameters, especially at the lower frequencies where these stimuli were perceptually wigglier. The surprising result was that the orientation discrimination thresholds were the same for the wiggly and the straight objects. Further both thresholds were the same as those obtained with classical orientation discrimination stimuli: straight-edged bars and two small dots. We inferred that object orientation can be determined in at least two ways: end-point locations can be compared and the central axis can be used. The observer appears to use the axis when it is straight, or nearly so, but not when it is highly modulated and thus not a good source of orientation information. We found no evidence that edge orientation is used directly. In summary, we found further evidence for a medial representation and strong evidence against existing edge-orientation-based models of perceived orientation (Paradiso, 1988). The results are reported the paper "Across-Object Relationships in Perceived Object Orientation" by Burbeck and Zauberman (in press) and were reported at ARVO 1995, "Perceived Object Orientation: Edges, Ends or Middles?"

4) Other Support for a Medial Representation

Our experimental findings that a medial representation does exist (Burbeck, et al., 1996; Burbeck & Zauberman, in press) have been bolstered strikingly by recent results from two other laboratories. Kovacs and Julesz reported (Kovacs & Julesz, 1994) that contrast discrimination sensitivity is enhanced at the medial locus, and Lee, Mumford, and Schiller (Lee, Mumford, & Schiller, 1995) reported finding neurons that respond when their classical receptive field contains the medial locus (and does not contain the boundaries of the object). Our core model appears to be a particularly timely addition to the field.

5) Current Research on Cores

(This work is being conducted under other sponsorship, but the results arose from research sponsored by this grant.)

Following Kovacs and Julesz's dramatic finding that contrast-discrimination sensitivity is enhanced at core locations, we began investigating another possible property of this special location: enhanced position sensitivity. Thus far we have data on a circle of a single size, and the results there are promising. In this experiment, we show a line drawing of a circle within which a small probe dot is located. We then show the line drawing again with the dot in a slightly different location. (We allow sufficient time between the intervals to eliminate apparent motion, and we add random jitter to the location of the circle on the display.) The observer's task is to report the direction in which the second dot is displaced relative to the first. The probe dot is always on a 45° line through the center of the circle, so the judgment is down and left or up and right.

If the observer were only using the shortest distance to the circle boundary as his location cue, then the position discrimination threshold should increase with increasing distance from that boundary, following Weber's law for size, and should be maximum at the center of the circle. The results show that this does not hold. Instead, position discrimination thresholds are highest at locations between the center and the edge and have a clear local minimum at the circle's center. The pattern of results looks much like Kovacs and Julesz's findings for contrast discrimination. We also investigated the effect of circle size to determine whether the area of enhanced sensitivity at the center of the circle scales with circle size, as predicted by the core theory, and found that it increases but does not exactly scale. Xiaofei Wang, a Psychology graduate student, completed these studies and tested the core theory prediction quantitatively. This work forms the primary part of Mr. Wang's master's thesis. A manuscript is in preparation.

B. From Objects to Sets of Objects

The shape representation process that we have modeled as core-formation and have studied using judgments of individual sizes and specific locations, results in a high resolution representation of a specific figural region. But an ordinary visual scene contains many regions that could be seen as figures, and they are not all represented with such resolution. How is spatial information about more complex scenes encoded? Our first step in working on this more general problem in object perception was to consider sets of similar objects. Such stimuli abound in both natural and man-made environments: a bunch of bananas, the flowers on a bush, the branches of a tree, the legs of a chair, desks in a classroom, cupboards in a kitchen, tiles on the floor, books on a shelf; the list seems endless. How are such sets represented?

We have obtained striking results on this topic which suggest a radically different way of thinking about the representation of spatial information. We have found that when a set of similar items is presented, the visual system summarizes the properties of the set quite precisely but retains no information about the individual items. This contrasts starkly with the common assumption that items in sets are represented individually at reduced resolution.

Our experimental approach was to measure directly the information that an observer has about a set of objects. Specifically, we measured what information the observer retains about individual objects in the set, and we measured what he retains about the characteristics of the set as a whole. In most of the studies, we used sets of spots of different sizes — size being a simple parameter with which we have considerable experience. Experiments were also conducted with sets of oriented line segments to ensure the generality of the results.

1) Membership

To determine what an observer knows about the individual items in a set, we conducted "membership" experiments. The observer was presented with a set of objects, as shown in Fig. 5. After viewing the set for 500 ms, the observer was shown a pair of targets, one of which was a member of the set he had just viewed, and one of which was not. (Non-members had sizes equidistant between member sizes.) The observer's task was to report which of the two test spots was a member of the previously viewed set.

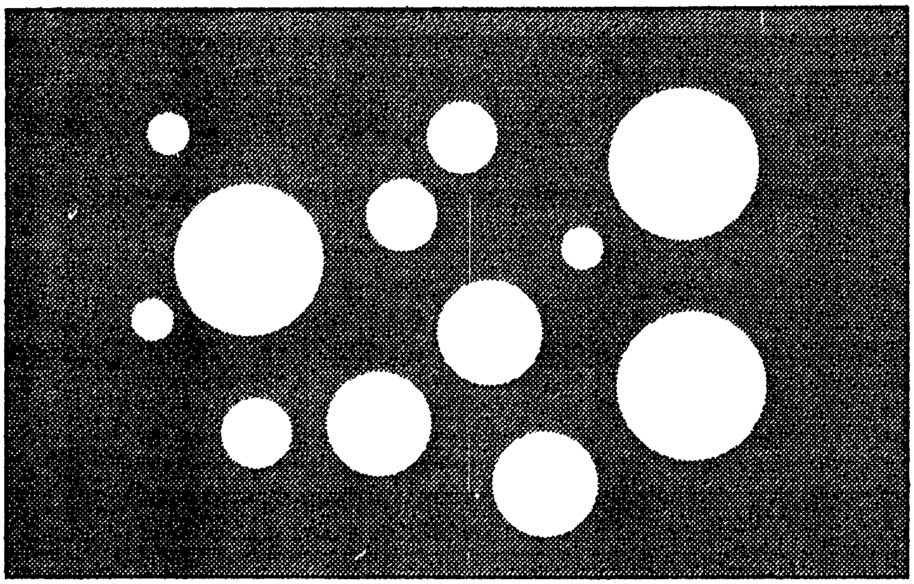

Fig. 5 An example of the stimuli used in the experiments on perception of multiple objects. The number of spots was varied; the average density of the spots was held constant.

Sample data are shown in Fig. 6. Performance on this task was at chance: observers chose members and non-members equally often even though the difference between the non-member and any member in the set was a highly discriminable 20% or more. (The standard size discrimination threshold is 3-5%.) Not shown in this figure is our finding that observers were much more likely to select the item in the pair that was nearer the mean of the set. Every item was paired with a larger and a smaller item, so this preference did not affect the average performance for a given item reported here.

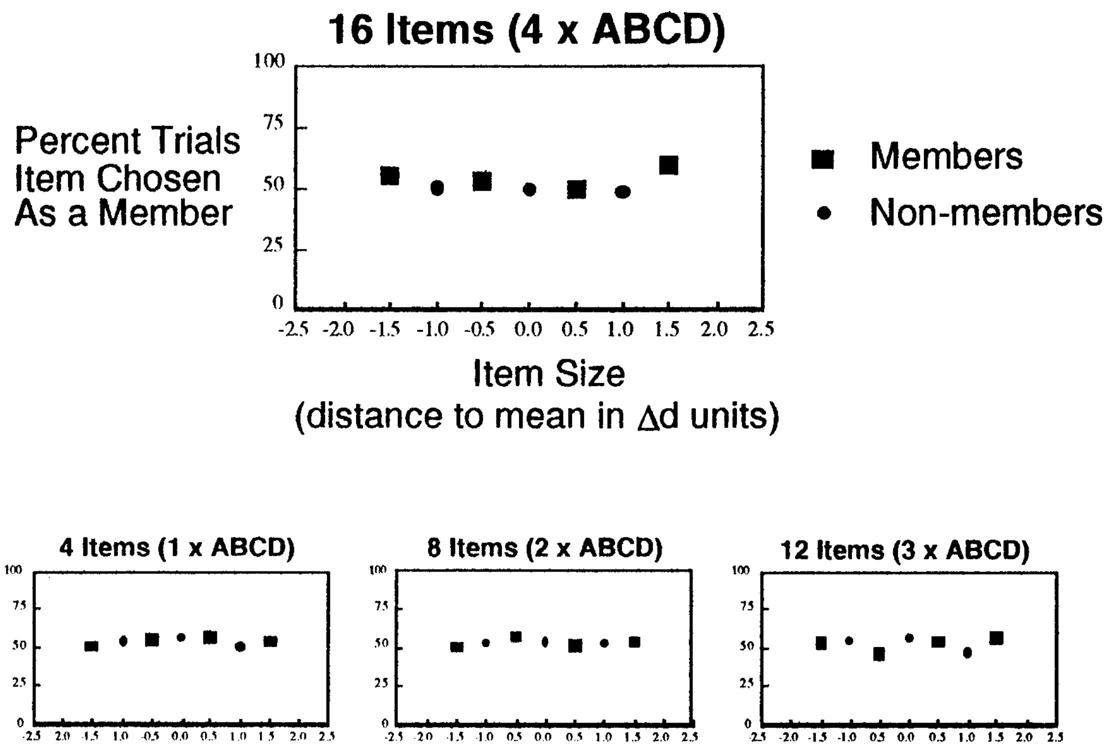

Fig. 6 Sample data from the membership experiment, showing observer's inability to distinguish members from non-members in a set of spots in which the spot sizes were incremented in 40% steps.

This preference for choosing items near the mean was also revealed in a yes/no experiment that we conducted. In this experiment, a set was shown and then a single test spot was presented. The observer was asked to indicate whether or not the test spot was a member of the set. Sample results are shown in Fig. 7. Observers showed the same insensitivity to actual membership, and showed strong sensitivity to the nearness of the test to the mean size of the set.

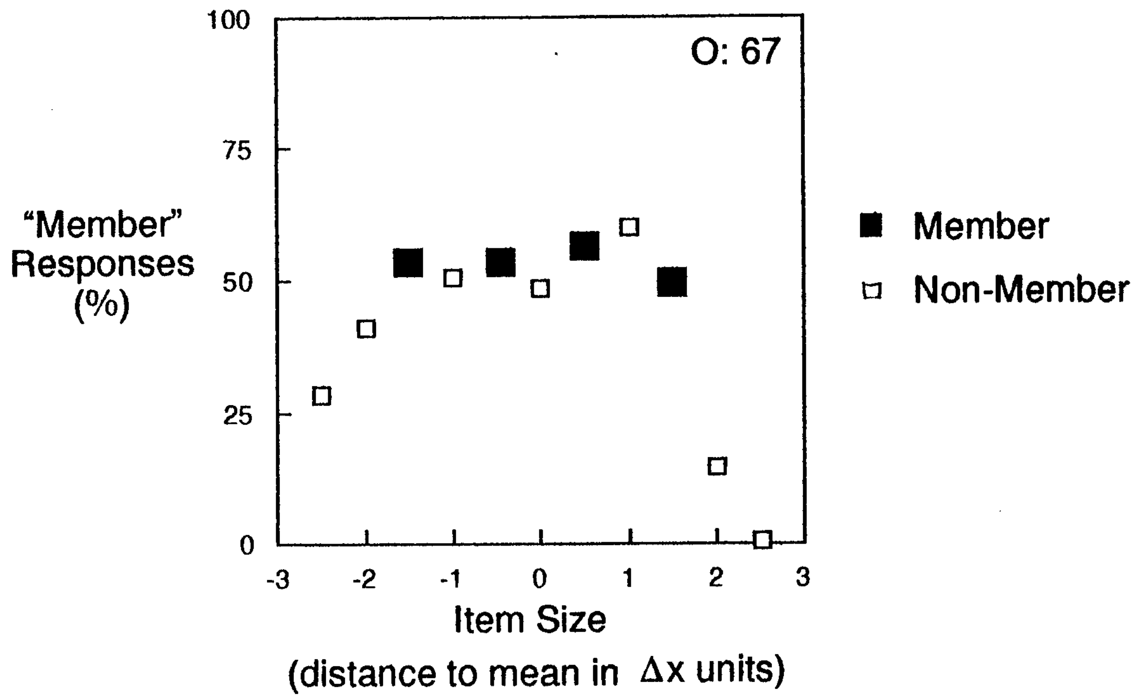

Fig. 7 Results of yes/no membership experiment measuring observer's subjective judgment of whether a given item was a member of the set. (Data shown are for the 40% distribution.)

Using the yes/no membership task, we have also found insensitivity to the form of the distribution — with only modest changes in the yes/no membership function between triangular and rectangular distributions.

Test spots for the set shown in Fig. 5 are shown in Fig. 8. The introspective reader can judge for himself his confidence that he knows which is the member. Our data show clearly that this information is not available.

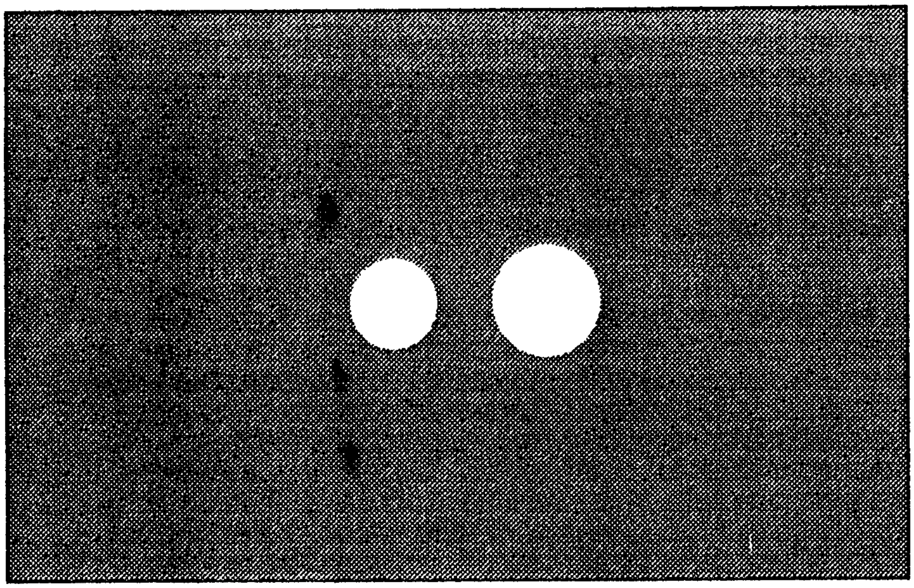

Fig. 8 Sample test spots from the membership experiment. One spot was a member of the set; the other differed in size from each spot in the set by 20% or more. 20% is several times the size discrimination threshold for single items.

In these membership experiments, it is the multiplicity of items, not the brevity of access to the stimulus, that limits performance. Although one might call the processing pre-attentive, it is significant that attention can be employed without improving performance. (Performance can be improved by presenting the test stimulus first; the task then becomes visual search.) The insensitivity of these results to scrutiny emphasizes the fact that observers are creating a representation of the set, not of the individual items. If the target item is cued first, as in a visual search task, then other types of processing can be employed. Our research focused on the perception of the set as such.

2) Mean

In the other type of experiment that we conducted, the "mean" experiment, we measured the observer's knowledge of the properties of the set as a whole. The observer was shown a set and then shown a single test spot. His task was to report whether the test was larger or smaller than the mean of the set (in a method of constant stimuli paradigm). Probit analysis of the results yielded the discrimination threshold for judging the mean size (i.e., the slope of the resulting psychometric function).

Mean discrimination thresholds varied from 3 to 12% depending on the distribution of sizes used. (There were also differences in sensitivity between observers.) The lowest mean discrimination thresholds were obtained with triangular distributions of 1, 4, or 9 spots and spot sizes separated by 5% increments. These distributions contained the mean size as one, two, or three of the spots (for the 1, 4 and 9 spot set sizes); the mean-sized spots could have been identified by the observer and used to support his judgment of the mean. To prevent this, a different type of distribution was used. We measured mean discrimination thresholds with uniform distributions of 2, 4, 8, 12 and 16 spots of constant range which did not contain the mean value as one of the sizes. (See Appendix F for further details of the stimuli.) Thresholds were 4-6% for sets in which adjacent sizes differed by 5% (yielding a largest size/smallest size ratio of 1.17) and were 6-12% (depending on the observer) for sets in which adjacent sizes differed by 40% (yielding a largest/smallest ratio of 2.74, as shown in Fig. 5). The increase in threshold was modest given the large increase in the range of sizes. For these rectangular distributions, the mean discrimination threshold was constant or decreased with increasing number of spots, suggesting (but not requiring) the operation of a parallel process.

3) Theory of Statistical Representations

We infer from the above results that multiple similar items are represented in a qualitatively different way than are individual objects, that the representation consists of precise (high resolution) information about average object properties, some information about the distribution of values present (as indicated by the membership data), and no information about individual objects in the set. We theorize that this "statistical" representation is associated with all items in the set, so that a region is represented as containing things of about this size, this average orientation, color, texture, etc., in this spatial density. The statistical representation does not assign specific parameter values to specific objects. This idea dovetails nicely with Treisman's ideas on the need for attention to bind features together (Treisman, 1988).

The statistical representation is an intriguing subject of study in its own right: it is a basic (and economical) form of spatial encoding with broad application. It may also play a role in guiding the perception of individual figures — a subject of interest to us in our work on cores. The characteristics of this set-encoding process may also provide an explanation for set-size effects, i.e., the reduction in discriminability that occurs when an item is embedded in multiple copies of another item, for visual pop-out (visual saliency in the presence of multiple other items), for the difficulty of some conjunctive searches, and for illusory conjunctions (which are predicted directly from the assignment of a set of properties to all objects in a region). Finally, set perception appears to be strongly related to the creation of visual categories over time. In previous studies (Burbeck & Swift, 1988), we found that observers could create very accurate representations of the mean size of a set of sequentially presented targets — the temporal equivalent of the spatial statistic.

A manuscript reporting all of our experimental results on sets is in preparation; they were also reported at ARVO in 1995 and 1996.

C. Area Perception

In related work, begun under sponsorship of Pizer's NIH grant, we conducted experiments aimed at the question: does the visual system encode spatial area directly or is it only inferred from length judgments. Several experiments were conducted, all of which point to the adequacy of length judgments in explaining "area perception". There were two major results. First, we found that within-shape area comparisons had the same precision as comparable length judgments whereas between-shape area comparisons had very elevated discrimination thresholds. Observers could not make accurate comparisons of the areas of differently shaped objects. The second result was obtained using a novel paradigm, from which one of our set measurement experiments was derived. In this paradigm, the observer was shown a pair of spots of different sizes together with a third spot, which was the test spot. The observer's task was to report whether the test spot was greater or less than the mean size of the two other spots. The observers responses were consistent with judgments of the average of the diameters of the spots and not consistent with judgments of the average of their areas. Some of these results were reported at the annual Psychonomic Society meeting (Burbeck & Ariely, 1992).

D. Occlusion Edge Blur

Overlapping interests among several of us resulted in a study on what we believe to be a new cue to relative depth: occlusion edge blur. Occlusion edge blur is the blur or sharpness of an edge between two regions, which are arranged so that one can be seen as occluding the other. We found that the observer's percept of which of the two regions is nearer to him is affected by whether the shared edge is blurred or sharp. The results of this study were reported in the paper, "Occlusion Edge Blur: A Cue to

Relative Visual Depth" by Marshall, Burbeck, Ariely, Rolland, and Martin, (in press, Journal of the Optical Society of America A). Fig. 1 of that paper illustrates the effect.

E) Multi-scale probe for measuring perceived 3D shape

The PI supervised a computer science graduate student's dissertation research on multi-scale measurement of perceived 3D shape. The student, Peter Brown, has developed a measurement probe, similar in spirit to the one used by Koenderink. As with the Koenderink probe, the observer adjusts the slant and tilt of a "top" (a disk with a normal stick) so that it matches the slant and tilt of the surface under it. Peter uses probes of various sizes and has the observer adjust the probe so that it lies parallel to the region under it — ignoring the smaller scale bumps and dips. To avoid the bi-tangency problem, i.e., the problem of having the slant and tilt of the probe determined by the two highest peaks under it, the probe is elevated above the surface a distance that is proportional to the size of the probe. Data have been obtained that show the desired scale-specific judgments. Peter's intent is to examine the effects of particular depth cues — texture and head-tracking — on the perception of 3D shape to determine if their effects vary with the spatial scale of the judgment. The overall goal is to provide a tool for assessing the value of particular 3D cues as a function of the scale of the critical information in the image.

Peter received a small amount of support from this grant in the prior year. The PI's time (roughly an hour a week) has been covered by this grant. The relation to the PI's main line of work is in the relationship to Weber's law and to the core theory which postulates that one core (roughly one spatial scale) is accessed at a time.

II. Publications, Presentations and Participants

A. Written Publications (Published or In Press)

Burbeck, Christina A., "Separation Discrimination with Embedded Targets", *Vision Research*, 32, 2295-2302, 1992.

Burbeck, Christina A. and Scott Hadden, "Scaled Position Integration Areas: Accounting for Weber's Law for Separation", *Journal of the Optical Society of America A.*, 10, 5-15, 1993.

Pizer, Stephen M., Christina A. Burbeck and Daniel S. Fritsch. "Human Perception and Computer Image Analysis of Objects in Images", *Proc. Conference of the Australia Pattern Recognition Society* (DICTA), Vol. I, pp. 19-26, 1993.

Pizer, Stephen M., Christina A. Burbeck, James M. Coggins, Daniel S. Fritsch, and Bryan S. Morse. "Object Shape before Boundary Shape: Scale-space Medial Axes", presented at *Shape in Picture*, (NATO Advanced Research Workshop), 1992. *Journal of Mathematical Imaging and Vision*, 4:303-313, 1994.

Pizer, Stephen M. and Christina A. Burbeck. "Cores as the Basis for Object Vision in Medical Images", *Proc. SPIE Medical Imaging '94: Image Perception*, Vol. 2166, pp. 191-198, 1994.

Morse, Bryan S., Stephen M. Pizer and Christina A. Burbeck. "General Shape and Specific Detail: Context-dependent Use of Scale in Determining Visual Form". *Proc. of the Second International Workshop on Visual Form*, pp. 374-383, 1994, World Scientific, Singapore.

Burbeck, Christina A. and Stephen M. Pizer. "Object Representation by Cores: Identifying and Representing Primitive Spatial Regions". *Vision Research* 35, pp. 1917-1930, 1995.

Burbeck, Christina A., Stephen M. Pizer, Bryan S. Morse, Dan Ariely, Gal S. Zauberman, and Jannick Rolland. "Linking Object Boundaries at Scale: A Common Mechanism for Size and Shape Judgments". *Vision Research.* 36, 361-372,1996.

Marshall, Jonathan A., Christina A. Burbeck, Dan Ariely, Jannick P. Rolland, and Kevin E. Martin, "Occlusion Edge Blur: A Cue to Relative Visual Depth". In press, *Journal of the Optical Society of America A.*

Burbeck, Christina A. and Gal Zauberman. "Across-Object Relationships in Perceived Object Orientation". In Press, *Vision Research..*

B. Manuscripts in Preparation

Burbeck, Christina A. and Dan Ariely. "One, two, three, many: Representing multiple similar objects". To be submitted to *Vision Research.*

C. Papers Presented at Meetings and Colloquia

Rolland, J. P. and Burbeck, C. A. (1992) "Depth and Size Perception in Virtual Reality Systems", Invited Address, Optical Society of America Annual Meeting, Albuquerque, September.

Burbeck, C. A. "Defining Perceptual Objects", Optical Society of America Annual Meeting, Albuquerque, September, 1992.

Burbeck, C. A. and Ariely, D. "Perceived Area: Within-Shape And Across-Shape Discriminations", Psychonomic Society, St. Louis, November, 1992.

Burbeck, C. A. "On The Road To Visual Object Perception: Linking Visual Regions", invited presentation, Psychology Department, NC State, November, 1992

[During 1993, the PI took a (mostly self-supported) sabbatical. No additional funds were allocated and research continued at a very reduced level. The second year of funding was thus 1994. The PI's goal was to delve into the theory more deeply without the daily interruptions of laboratory work. It was during this contemplative year that the core theory reached fruition.]

Burbeck, C.A. and S. M. Pizer, "Object Representation by Cores", ARVO annual meeting, 1994.

Burbeck, Christina A. and Stephen M. Pizer, "The Scale of Shape Perception", Psychonomic Society Annual Meeting, 1994.

Burbeck, Christina A. and Stephen M. Pizer, "Figure Formation In Human Vision: Psychophysical Results And A Model". Invited talk delivered by S.M. Pizer at Koenderink's laboratory in Rijksuniversiteit Utrecht, September, 1994.

Morse, B. S., S. M. Pizer, and C. A. Burbeck, "General Shape and Specific Detail: Context-dependent Use of Scale in Determining Visual Form", Presented at *International Workshop on Visual Form*, Capri, 1994.

Pizer, S.M., and Burbeck, C.A. "Cores As The Basis For Object Vision In Medical Images", *Medical Imaging '94: Image Perception*, SPIE, 2166:191-198, 1994.

Burbeck, C.A. , Participant in Axiomatic Vision Workshop, Computer Science Dept., UNC, Chapel Hill, Nov. 9-11, 1994.

Burbeck, C.A., "One for the tree and one for the forest: Two types of spatial representation in human vision. " Colloquium, Mar. 1995, Computer Science Dept., UNC.

Ariely, D. and Burbeck, C.A. "Statistical encoding of multiple stimuli: a theory of distributed representation." ARVO annual meeting, 1995.

Burbeck, C.A. and G. Zauberman, "Perceived object orientation: Edges, ends or middles?", ARVO annual meeting, 1995.

Burbeck, C. A. and Wang, X. F., "Enhanced position sensitivity at scaled medial loci", ARVO annual meeting, 1996.

Burbeck, C.A. and Maher, S. J., "Identical distractors aren't identical", ARVO annual meeting, 1996.

Pizer, S. M., and Burbeck, C. A., "Robust object representation", ARVO annual meeting, 1996.

D. Consulting and Advising (by the PI)

Member of NIH, Visual Sciences B study section Feb. 1993 - Feb. 1996.

Member of NSF Sensory Systems advisory panel, Sept. 1989 - Sept. 1992.

Reviewer for many journals and other funding agencies.

E. Professional Personnel Associated with Research Effort

1) Investigators funded by this project:

Christina A. Burbeck, Research Professor of Computer Science and Psychology

Dan Ariely, Graduate Student in Psychology.

Bryan Morse, Graduate Student in Computer Science, now Assistant Professor of Computer Science, Brigham Young University.

Peter Brown, Graduate Student in Computer Science.

Sean Maher, Graduate Student in Computer Science.

Gal Zauberman, undergraduate Research Assistant.

Shelley Poovey, undergraduate Research Assistant.

Michael Capps, undergraduate programmer.

Scott Hadden, Research Assistant

Irene Snyder, undergraduate Research Assistant.

2) Collaborators who have worked on the above research, but were not funded by this grant:

Stephen M. Pizer, Kenan Professor of Computer Science.

Jonathan Marshall, Assistant Professor of Computer Science.

Jannick Rolland, Assistant Research Professor of Computer Science.